Place with computer magazines

11>

0 74470 92596 5

Issue #15 Nov/Dec 1999
$9.95 US $14.95 Canada

Hollywood on your Desktop...

Storm of the Century images courtesy of Netter Digital Entertainment, Inc. Storm effects added with Particle Storm 2 Pro.

The Hunley image courtesy of Station X Studios, LLC. Explosion created with Particle Storm 2 Pro; water created with NatureFX.

AVAILABLE FROM

800-544-6599 • 414-548-8120
WWW.SHARBOR.COM

The most popular LightWave plug-in in the world!

Advanced features include: Particle Fields, Hair, Spawner, Sub-Objects, Sub-Polys, Mimic, Assembler, Envelopes, Layout Integration, Sticky Particles, Bounding Shapes, Collision Detection, Collision Avoidance, OpenGL Realtime Preview, Gravity, Wind, Flocking & Swarming, Drag, Force Fields, Death Wish, Particle Shader, and Parenting!

Realism Unparalleled...

The software responsible for brining the Titanic back to life is now available for LightWave! NatureFX includes these advanced features: NatureFX Clouds for creating 3D volumetric clouds, NatureFX Air to simulate the coloration of the sun, horizon, and sky, and to generate haze and rainbows. NatureFX Ocean creates waves, wakes and ripples, and NatureFX Object Bob gives you bobbing objects.

800-944-5520 • 702-990-2277
WWW.DYNAMIC-REALITIES.COM

The Editor
Doug Kelly *dakelly@megalink.net*

Advertising Rep
Tiffani Banaszak *keyframe1@aol.com*

Art Director
Tiffani Banaszak *keyframe1@aol.com*

Co-Conspirators, Writers, and Artists
Sam J. Bowling *sbowling@home.com*
Jason Goodman *jason@msd-nyc.com*
Julian Kain *julian@juliankain.com*
Graham McKenna
graham@grahamprismatic.free-online.co.uk
Brenden Mecleary *bman2@ix.netcom.com*
Matt Morgan *violator@thethreshold.net*
Steve Pugh *steve@foundation-i.com*
Charlie Ramos *charlieramos@hotmail.com*
Lee Stranahan *lee@stranahan.com*
Phil South *phil_south@hotmail.com*
David Warner *dwarner@albany.net*

The Digital Underground
DMG Publishing
2756 N. Green Valley Pkwy. #261
Henderson, NV 89015
702-990-8656

How To Get It
888-778-WAVE
http://www.keyframemag.com

Keyframe Magazine is published 6 times annually by DMG Publishing. © 1999 DMG Publishing. All rights reserved. Reproduction, in part or in whole, without written permission is strictly prohibited. Printed in the USA.

The Secret Formula
Keyframe Magazine is created on Intergraph machines using Quark 4.0, Photoshop 4.0, LightWave 5.6, Painter 5, and various plug-ins. Hardware includes a MusTek 2400 9600dpi scanner, Iomega Zip and Jaz drives, and a DPS Perception. The cover CD is created by David Warner using a GateWay2000 PPro200 w/128 mb Ram and an 18GB hard drive, a Truevision Bravado 1000 video capture card, Sony CD-Recorder, and an Iomega Zip drive. Software includes LightWave 5.6, Aura, Photoshop 4.0, Macromedia Director 6.5, Adobe Premiere 4.2, QuickTime 3Pro, and Adaptec Easy CD Creator.

No dongles were harmed in the making of this publication.

Departments

3 Editorial — by Doug Kelly
- Be Careful What You Wish For!

4 Ear to the Wall
- DPS' new Velocity
- Image Entertainment, Inc. to release Odyssey tapes on DVD
- Safe Harbor's two new web sites
- Digital Hollywood at the 2000 International CES
- The mailing list's new digs

10 Macintosh Up Close and Personal — by Jason Goodman
- Comparing the ToasterNT to Digisuite LE

12 On Topic — by David Warner
- 64K limit, Morph remodel, Skinning, Standby

31 Take One
- Onscreen Effects

34 Reviews
- Surfaces ROM
- Dosch Textures vol. 1
- GeForce 256

37 Final Cut — by Steve Pugh
- Foundation Imaging's Workflow Pipeline for RSST

Tutorials

6 Character Shop — by Doug Kelly
- Animating a manual match move

14 Programmer's Hideaway — by Brenden Mecleary
- Drawing bitmaps in LWPanels

16 Metamorphosis — by Charlie Ramos
- Case study of a short film

18 Made to Measure — by Graham McKenna
- Photorealistic leather texture mapping

22 Getting Started with LW Modeler — by Sam J. Bowling
- Before you start modeling

26 Lighting to Match Plate Photography — by Matt Morgan
- Matching real-world lighting

29 Radio-sity Music Hall — by Phil South
- Simulating reflected light

EDITORIAL

Be Careful What You Wish For

I ended last issue's editorial with the wish, "May we all continue to live in interesting times!" Let that be a lesson to me to be more careful about what I wish for, because the past two months have been more interesting than I *ever* care to experience again.

As I requested, some of you contacted me with feedback on the past year of *Keyframe*, and what you want to read here in the future. The overwhelming majority of that feedback has been positive, and it is particularly gratifying that the highest-ranking requests were for changes that are already evident in this issue. We are continuing to publish as many tutorials as possible. In order to pack this issue's *seven* tutorials into the page count dictated by our budget, we have dropped our usual 10-point font to 9-point. This makes the magazine's page count smaller, but you are actually getting *more* content!

One change came only after long consideration: we are dropping the CD-ROM entirely after this issue. Repeated and varied failures by more than one CD duplication service have delayed subscribers' issues far too long. Unique and original content on the CD has fallen as low as 5MB in some issues, with most of the space filled by materials already available elsewhere. In weighing the options, we decided that it would better meet all our readers' needs if we divert resources from CD-ROM production to upgrading and maintaining the *Keyframe* website. This will enable us to get *Keyframe* to you more quickly, provide faster feedback for both subscriptions and editorial support, and provide most of you with a more timely and efficient means of accessing the complementary digital files. We will not charge for online access, and current CD subscribers' term will be extended two extra issues for no charge. I want to make very clear that we are not changing *Keyframe* to an online magazine format. We will not go there, not as long as I'm the Editor!

I'm apparently not the only one who has been living in interesting times. The LightWave Mailing List switched over to eGroups in mid-November, the most drastic change the list has seen in a long time. Let's all work to preserve the essentially helpful and friendly tone of the old list. Whatever the outcome, I think we all owe David Warner a huge "Thank You!" for all the time and effort he has donated to maintaining the LWML. As of this writing, several LightWave-related products are still nearly ready to ship, but not ready enough to make reviews for this issue. On the positive side, the developments I have been watching in L[6], Kaydara's FiLMBOX, and Station X Studios' Project: messiah definitely indicate that these tools will be worth the wait. By the time you read this, I believe most or all of these products will be shipping, and you can look forward to in-depth reviews and tutorials in the next issue. In the meantime, this issue of *Keyframe* brings you tutorials on lighting, surfacing, modeling and plug-in coding that will apply to nearly any version of LightWave.

Happy holidays, and an especially Happy New Year!

Doug Kelly

Editor

Ear to the Wall

News from the LightWave arena

What Velocity do you move at?

dpsVelocity(tm) is a new real time non-linear editing system for digital video professionals. A joint effort of DPS' award-winning hardware and software development teams, dpsVelocity delivers faster, sharper and easier edits while raising the bar for reliability and affordability.

At the heart of every dpsVelocity editing system is the new dpsReality studio digital disk recorder card. According to DPS, this advanced hardware, which supports compressed and uncompressed video recording modes, delivers quality so sharp it surpasses Digital Betacam. Amazingly compact, the dpsReality card combines the functions of a dual stream DDR, video mixer/keyer, real time 3D effects processor, multi-channel audio mixer, SCSI controller and two graphics framestores in a single PCI slot.

It takes a powerful software application to exploit the capabilities of the dpsReality hardware. The extensive real time capabilities of dpsVelocity can do just that. Optimized for the way you work, the intuitive and flexible interface makes it easy to perform even the most complex three and four point edits. Numerous real time filters such as color correction, speed change and transparency are included along with hundreds of ready to use 2D/3D transitions, rolling/crawling titles and perspective, rotation and warp effects. Enhanced dpsVelocity audio features provide real time volume, pan and EQ control for eight channels (four stereo pairs) of synchronized audio.

For more information, visit www.dps.com.

Opening the Mind's Eye

CHATSWORTH, CA (August 18, 1999) - Image Entertainment, Inc. (NASDAQ: DISK), a leading licensee and distributor of DVDs in North America, today announced the signing of an exclusive agreement with Steven Churchill's Odyssey Productions for the release of six state-of-the-art computer animation programs on DVD and videocassette.

Under the terms of the agreement, the first title to be released will be *Computer Animation Marvels*, a collection of 17 animated shorts from four continents that defines the current state-of-the-art for this celebrated programming. Culled from the world's top CGI artists and storytellers, *Computer Animation Marvels* includes the shorts "Bingo," an award winning look at identity crisis, "Pets," which deals with the innermost secrets of our furry friends, "The Physics of Cartoons," and 14 additional humorous and thought-provoking films.

Odyssey Visual Design was founded by Steven Churchill in 1982 and was the first San Diego company to produce 3-D computer animation. In 1988, Odyssey produced the world's first entertainment-oriented computer animation compilation for home video release, *State of the Art Computer Animation*. The release won the "Best Video Art" award from the American Film Institute that year. In 1990, the company co-produced the landmark video, *The Mind's Eye: A Computer Animation Odyssey*. Over 400,000 copies of the video have been sold to date and it rode the Billboard video sales charts for over 30 weeks. Odyssey followed this success with *Beyond the Mind's Eye, Computer Animation Festival Volume 1, Virtual Nature, Imaginaria, Computer Animation Festival Volume 2* and *The Gate to the Mind's Eye*.

In 1995, Churchill renamed his business Odyssey Productions which has produced *Odyssey Into The Mind's Eye, Computer Animation Festival Volume 3, Turbulence, Computer Animation Classics, Cyberspace, Computer Animation Showcase, Luminous Visions, Ancient Alien* and most recently *Computer Animation Celebration*.

Visit Image on-line at http://www.image-entertainment.com for corporate and marketing information, DTS information and the latest in DVD, laserdisc and Image-exclusive VHS releases. Also visit Ken Crane's DVD and Laserdisc Superstore on-line at http://www.kencranes.com for consumer direct sales.

Safe Harbor Founds Two Innovative Web Sites

VideoToasterNT.com and New JoyOfSix.com Offer Users Product Information, Instant System Quotes and More!

Waukesha, WI, November 19, 1999-- Safe Harbor Computers, Inc. is proud to announce two innovative web sites to provide information about the hottest NewTek products. VideoToasterNT.com and the new JoyOfSix.com were developed by Safe Harbor to promote a technological community made possible by the Internet. These new sites allow for Video Toaster NT and LightWave [6] current and potential users to come together and share tips and information, as well as get expert, accurate advice about these products.

VideoToasterNT.com
The Video Toaster NT web site (http://www.videotoasternt.com) was designed to provide up-to-date, accurate information about the Video Toaster and to provide detailed product specifications. The site's interactive forum allows for questions to be answered by NewTek and Safe Harbor representatives as well as other Video Toaster NT users. Designed for both current users of the Video Toaster NT as well as those considering it, the site provides a frequently-asked-questions page and related links to other Video Toaster NT resources.

Updated daily, this site is also valuable for technical resources. Available on-line are technical support documents for the Video Toaster NT, and tips for getting your system up and running. In addition, this site offers instant system quotes on custom turn-key complete Video Toaster NT systems from Safe Harbor and easy click ordering. The Video Toaster NT can't be made any easier!

JoyOfSix.com
The newest web-site offered by Safe Harbor, the JoyOfSix web site (http://www.joyofsix.com) was created for individuals whose knowledge of NewTek's LightWave [6] ranges from novice to expert. The site provides detailed information about LightWave Version 6 for those interested in learning more about this incredible product. JoyOfSix.com also allows you to custom design the LightWave animation workstation of your dreams and provides instant system quotes from Safe Harbor. Still in development, JoyOfSix.com will feature a great deal more in the future including an artists' gallery where LightWave animators can showcase their work for other animators.

About Safe Harbor Computers
Established in Waukesha, Wisconsin, in 1987, Safe Harbor Computers is home to graphics and video products (hardware and software) for Windows, Mac, DEC Alpha, and Amiga computers. With friendly, knowledgeable sales and technical staff in additional to highly-competitive prices, Safe Harbor continues to set the standards for the industry. Safe Harbor's award-winning, graphically advanced web sites operate with the latest technologies, such as secure, one-click online ordering, automatic confirmation of orders, order tracking, instant quotes, user reviews, and much more. Visit Safe Harbor online at http://www.sharbor.com or call 1-800-544-6599 for more information about products and services.

Hollywood comes to Las Vegas

The 2000 International CES, to be held at the Sands Convention Center in Las Vegas January 6-9, 2000, has announced the additional of a special forum: Digital Hollywood. Digital Hollywood includes special, targeted conference sessions and LVCC exhibit space to showcase the digital marketplace of the future.

Conference Sessions
Sessions focus on views of the digital home, convergence strategies and realities, E-commerce, investment and venture capital in entertainment and technology start-ups and more. Since 1990, more than 15,000 industry executives have attended Digital Hollywood.

Some of the conference session topics include *The TV and PC Internet Connection - Technology, Commerce & Content*, *Digital Entertainment Convergence: A Content & Technology Explosion - TV, Music Cable, Internet, Update on Digital Piracy & Copyright Infringement,* and *Interactive TV/DTV Software & Programming Solutions*.

Digital Hollywood conference registration includes access to all exhibits, keynotes, super sessions, and power panels. Early registration for a full pass is $235, a daily pass is $150. On-site registration costs $300 for a full pass and $175 for a daily pass.

More information can be found at www.cesweb.org.

LightWave Mailing List has New Home!

As of November 15th, 1999 the LightWave Mailing List has officially moved to a new server based on the eGroups.com site.

There are two ways to subscribe to the new list: the quick-and-easy way that does not allow you to take advantage of any of the special subscriber options, and a slightly more complicated way that will give you full access to all of the cool features at eGroups.

Subscribing The "Quick & Easy" Way
If you want to subscribe to the new list and you want to receive individual messages without any bells & whistles, then all you need to do is send an e-mail to LW3D-SUBSCRIBE@EGROUPS.COM and a few minutes later you will receive a confirmation e-mail that you need to reply to in order to finalize your request to subscribe.

For information on the "Better Way", which offers many features such as a daily summary, visit www.tv3d.com/LWC.html.

Animating a Manual Match Move

One of the most demanding and popular uses of LightWave character animation is integrating a character into live-action footage. Doing this convincingly requires your character to interact accurately with the filmed environment, including lights, shadows, surfaces and objects. If the camera is not locked down, you will need to animate LightWave's camera to match the live-action camera's movement and lens. If the original camera was not a motion control rig (as described in *Keyframe #11*) you will have to perform a match move.

Your first goal in setting up a match move is to automate as many variables as possible. You have to perfectly coordinate the movement and rotation of the camera in all three axes, plus lens parameters. That's at least eight variables, any one of which can be changing on every frame. The more variables you can track, measure on set, or otherwise nail down in advance, the easier and more accurate your move match will be.

When you have automated as many variables as possible, prioritize the remaining ones in order of frequency of change, lowest first. The idea is to manually keyframe the lowest-frequency motions first. This eliminates one more variable with a minimum of effort, and makes keyframing the remaining higher-frequency variables easier by removing ambiguities between related variables. For example, an increase in the apparent size of an object could be the result of a Z-axis move (dolly in) or an increase in lens length (zoom), and is therefore ambiguous. If the lens is fixed or has already been animated according to the focus puller's information, you can automatically assume that the change is caused by a Z-axis move, and animate the CGI camera to match. You will also find that setting one variable's keyframes often reduces (or even eliminates) the need for keyframes for another variable.

One way to automate or eliminate a lot of keyframes is to stabilize the footage using compositing software like eyeon's Digital Fusion. One-point stabilization can remove vertical and horizontal shaking, but two-point stabilization can also correct for rotation and scaling. I ran the example footage through two-point stabilization with an oversized margin, so the edges of the images would not be lost in scaling or rotation. Figure 1 shows one frame of the original footage, the same frame with the safety margin added, and again after the two-point stabilization. The reference points for stabilization were the right and left ends of the door's threshold. This stabilization means those two points will be in exactly the same position within every frame of the footage.

The two-point stabilization still leaves a lot of movement in the footage; the farther from the reference points, the greater the motion. One way to remove a lot of this motion is to run a single-point tracker, and apply that tracker's motion to the LightWave camera. Digital Fusion enables you to export tracking data as a LightWave .MOT file, making this a fairly easy process. The Digital Fusion motion coordinates are changes, not absolute measures, and are relative to the default original coordinates of 0.5, 0.5, the center of the tracked image. You need to remember this when you apply the motion file to a LightWave object. The example footage, objects and motion file are included on this issue's CD-ROM.

Open LightWave and clear the scene. Set the Last Frame to 265. Set the Camera to zero on all axes of rotation and translation, and set a key on frame 0. Switch to Camera View. Open the Camera panel. The camcorder I used to record this footage has a 1.6 to 3.6 focal length zoom lens. I shot this at approximately 1.0 zoom factor, so set Lens Focal Length to 2.4 mm as a first approximation. Set Film Size to 1/2" CCD. Turn on Custom Size, and set Width to 2160 and Height to 1458, to match the resolution of the stabilized footage. Close the Camera panel.

Open the Images panel. Click Load Sequence and choose the Stable000.jpg sequence from the CD-ROM. Close the Images panel. Open the Effects panel. In the Compositing tab, select the image sequence as the Background Image. Close the Effects panel. Open the Options panel. In the Layout View tab, turn on BG Image for Layout Background. If you want to see the background in color, turn on OpenGL in Layout View Graphics. Close the panel. In frames 1 through 265, you should now see the stabilized footage; there is no frame 0 in the image sequence, so the background is blank on that frame.

Add a null and name it OriginNull. (First person to make a pun gets thirty whacks with a wet noodle!) Parent the Camera to this null. Load object 1_meter.lwo, located on the CD-ROM. This is your reference frame, a simple 1 meter square. Parent this object to the OriginNull. Move 1_meter to 2.0 meters on the Z axis, and set a key on frame 0. This 2.0 meter offset is not completely arbitrary; it's a rough estimate of the distance from camera to tracker pattern in frame 154. The 1_meter object is used strictly for reference during the match move, and should not appear in the final rendering. Open the Objects panel, and in the Appearance Options tab, turn on Unseen By Rays and Unseen By Camera and turn off all three

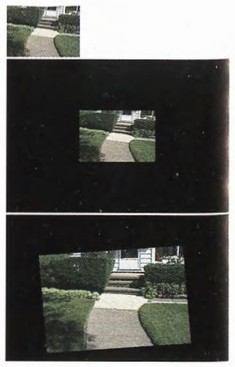

Figure 1: Top, one frame of the original footage. Middle, the same frame with safety margin added. Bottom, after the two-point stabilization.

shadow options. This will make the object invisible to the Camera during rendering. Close the Objects panel. You only need to see the reference frame as a wireframe object. Open the Scene Editor, and change 1_meter to wireframe display. Close the Scene Editor.

Add a null. Name it TrackerNull, and parent it to 1_meter. This is the object that will move to follow the tracker pattern. With TrackerNull still active, open the Graph Editor. Click Load Motion. Choose the TrackerNull.mot file from the CD-ROM. Click Use Motion. Your scene should look like Figure 2.

TrackerNull is not positioned anywhere near the second-step chip that I used as the tracking pattern. Tracking is based on XY plane movement of the tracker pattern. The motion file's XY coordinates are expressed as fractions of the image dimensions, that is, a pattern in the center of the screen would have the coordinates 0.5, 0.5. No value can be lower than 0.0, 0.0 and no value can be higher than 1.0, 1.0 because those coordinates are off the screen where the tracker could not follow the pattern. These coordinates are precise to six decimal places. The catch is that they are intended to match an image with the same proportions as the tracked footage. That means you need to resize the reference frame, the 1_meter object, to fit the boundaries of the background image as it appears in LightWave's Camera View. This is easy to tweak with the Numeric input. Select 1_meter, and change the Size until its edges precisely match the boundaries of the background image in the Camera View, as shown in Figure 3. Use the Numeric panel to tweak the Size in X and Y axes to at least three decimal places. Set the Z value equal to the X value. Take your time and get it right; this will affect every other step in the move match.

When you are satisfied with the Size, set a key at frame 0. If you step through the frames now, the TrackerNull should precisely match the chip in the second step.

Add another null, name it PorchPivotNull, and parent it to TrackerNull. Load the Porch.lwo object from the CD-ROM. Parent it to the PorchPivotNull. In the Scene Editor, set the Porch to full wireframe display. The nulls and reference frame can be bounding box wireframes, but you'll need to see all the Porch's edges in this project. With the Porch object selected, open the Surfaces panel. Select the PorchSurface surface, set its color to black, and click the T button next to Surface Color to open the Color Texture subpanel. For Texture Type, select the stabilized footage as a Front Projection Image Map. Set Texture Opacity to 100%, and close the subpanel to return to the Surfaces panel. This will make the Porch render as nearly invisible except for shadowing effects. In the Advanced Options tab, set Alpha Channel to Shadow Density. This will make the porch surface a shadow-catcher for any other objects you animate over it. Close the Surfaces panel. With the Porch object still selected, open the Objects panel. Turn off Self Shadow and Cast Shadow. In the Appearance Options tab, turn on Polygon Edges and set the Edge Color to bright blue. Close the Objects panel. This will make the wireframe outline of the Porch visible in your test renderings, making it much easier for you to evaluate the matching of edges to the background image. The critical details of the background image can be very small and difficult to see in some frames, and the wireframes in Layout can obscure those details. The rendered Polygon Edges are a very good visual aid. Just remember to turn off Polygon Edges before you start the final rendering!

Child objects inherit the scaling of their parent objects. We don't want the Porch object to be scaled to match the background image's aspect ratio, so we need to resize the PorchPivotNull to counteract the scaling of the 1_meter object. Resize the PorchPivotNull to the inverse of 1_meter's scaling. Figure 12.14 shows 1_meter object's scaling to be 6.345, 4.260, 6.345. Dividing 1 by each of these numbers gives a scaling correction of 0.1576, 0.2347, 0.1576. Apply these scaling corrections to PorchPivotNull. You can check the scaling correction, if you like, by loading another copy of the Porch object and positioning it to match the first one. You'll have to give it the vector sum XYZ translations and HPB rotations of all the parent objects to get a precise alignment for your comparisons. When you're done with the check, delete the duplicate Porch object.

The Porch object is constructed according to a sketch and measurements I made on location, with the model's pivot point set at the location of the chip in the second step that was used as a tracking pattern. This means that applying the XY motion of the tracker to this model will automatically make it match the XY motion of the real object in the background footage with no offsets. The stabilization of the footage, combined with the tracker-derived XY animation of the model, takes the place of two-thirds of the animation of the camera. The only positional variable left to animate is the Z-axis distance between the camera and the model. In order to animate the Camera, you would have to animate the 1_meter reference object's scaling to keep it exactly matched to the edges of the background image in the Camera view. You would also have to animate the counter-scaling of the PorchPivotNull to keep the tracker motion accurate. That's the reason for all the nulls and parenting; you can simply scale the OriginNull to adjust the distance between Camera and target, and the 1_meter reference

Figure 2: Tracker motion file applied to TrackerNull, which is parented to 1_meter reference frame. Note motion path trailing below null.

Figure 3: Resizing the reference frame, which is barely visible as a narrow border around the background image, scales the motion of TrackerNull to match the tracker pattern's position.

object automatically scales to match. You still have to counter-scale the PorchPivotNull, but since everything is proportional, you can do that with a single click-and-drag.

In the example footage, the camera moves in a complex path, partially orbiting the porch steps. This is fairly difficult to animate. However, the camera-and-porch relative motion can also be described as a comparatively simple set of rotations of the porch around the pivot located at the tracker pattern. That's why the Porch object is built with the pivot point in that location.

The camera's focal length determines the amount of parallax at any given distance. With the model located at the appropriate distance from the CG camera and positioned over the background image, you can adjust the CG camera's focal length to match the parallax of the CG model to that of the original object. If you have an accurate measure of either the camera-to-target distance or the camera's focal length, finding the missing measurement is not difficult. However, if you have neither the focal length nor the camera-to-target distance, you have to twiddle and tweak until you find a usable setting. This is why it is so important to record this information when you are on set, and why you should be especially nice to the focus puller. Fortunately for the purposes of this project, the initial setting of 2.4 mm is usable.

With this setup, you need to animate only a few, low-frequency variables to complete the move match: Size of the OriginNull to approximately match the distance from Camera to target, substituting for Camera Z-axis motion; Size of the PorchPivotNull in all three axes to control the height, width and depth of the Porch; and Rotation of the Porch in all three axes. As I mentioned previously, a useful strategy is to animate the lowest-frequency variable first. That means animating the OriginNull scaling first; you may only need a handful of keyframes, since you can be pretty sloppy with this variable and it doesn't change very quickly. The interaction of the PorchPivotNull scaling and the rotation of the Porch object mean that you will have to animate these variables pretty much simultaneously.

Go to frame 154. This was the reference frame for stabilizing the footage, so there should be no position, rotation or scaling distortions in this frame. The distance from the tracker pattern (the chip in the second step) to the camera in this frame is roughly two meters, so you don't need to scale the OriginNull. With no scaling of the OriginNull, you shouldn't have to make major scaling corrections with the PorchPivotNull. Rotate the Porch object so the line representing the upper front edge of the second step, which passes through the pivot point, is aligned with the upper front edge of the second step in the background image. Modify the Heading, Pitch and Bank values so as many as possible of the horizontal lines in the model align with horizontal references in the background image. Your results should look something like Figure 4. Set a key for the Porch at frame 154.

The Porch still appears a little too large. Tweak PorchPivotNull size. Try all three axes at once, first, then tweak individual axes to adjust width, height and length. You may find that you need to go back and tweak the Porch rotations to get a more precise alignment with the background. When you think you have matched the background's scale and rotations as closely as possible, set keyframes for the Porch, PorchPivotNull and OriginNull. You can check your work by rendering the frame with the same settings you will use for your final renders. With the Polygon Edges still turned on for the Porch object, your results should look something like Figure 5.

Repeat the above steps for the other key frames in a bracketing pattern. I suggest frames 1, 75, 180, 211, and so on. First, size the OriginNull. Examine the background image, and make your best estimate of the distance from the chip pattern to the camera. Divide that estimate by two (remember, 1.0 for the OriginNull is two meters' distance), and enter the result as the size of OriginNull. Second, rotate the Porch to align its horizontal lines with the corresponding horizontal edges in the background image. Third, size PorchPivotNull to bring the width, depth and height of the Porch into alignment with the background image. Finally, go back and tweak Porch rotations and PorchPivotNull size to correct any remaining errors. Set keys for all three elements, and go on to the next frame.

The fewer the keyframes, the smoother the motion and the less likely you will make errors that produce chattering. With every keyframe you set, the discrepancies to be corrected in the adjoining keyframes become smaller. It's up to you when to stop the tweaking; on a short deadline, 'good enough' may only be the six initial keyframes. With a longer production schedule, you may have time to tweak a couple of dozen keyframes.

For frames 212 through 265, you have a slightly harder job. Because the chip on the edge of the second step is no longer visible in the background image, you don't have Tracker data in the imported motion file. That means you don't know where the Porch is supposed to be located in the XY plane. That adds those two variable to the list you have already animated, and the interdependencies between all these variables can be a very demanding juggling act. This is where you practice your interpolation skills.

Your first step is to eyeball what variable settings you can simply by looking at the background image for frames 212 through 265. The most obvious candidate is the OriginNull scaling. Go to frame 212. In the Scene Editor, choose Hide All Objects. This gets the wireframes out of your way so you can see the background better. Scrub through frames 212 to 265, looking only at the distance between the camera and the porch. At what frames does the camera stop moving toward or away from the porch? Ignore side-to-side, pitch, bank or heading changes; concentrate on that Z-axis motion. By my estimate, the camera stopped moving closer to the porch on frame 231. Go to frame 231, or whichever frame you think is correct. Estimate the distance from the camera to the estimated position of the tracker pattern. You'll have to extrapolate the leading edge of the second step. Divide that estimate by two (remember, 1.0 for the OriginNull is two meters' distance), and enter the result as the size of OriginNull. Set the same value for a key on frames 232 and 265. That nails down the Z-axis motion for the rest of the match move. Another approach is to select the OriginNull and open the Graph Editor. Choose the X Scale, and look at the curve. You should be able to see a definite shape to the curve, and estimate pretty closely the value that will extend the curve to frame 231. Click Use Motion to close the Graph Editor. The scaling of the OriginNull directly affects the scaling corrections for the PorchPivotNull. Set keys at OriginNull's keyframes for the scaling of PorchPivotNull. As a first approximation, keep the same PorchPivotNull scaling as in frame 211.

The next lowest-frequency variable is the Porch object's Bank rotation. In the Scene Editor, make the Porch object visible as a full wireframe. On frames 231 and 265, adjust the Porch object's Bank rotation to align the horizontal edges of the Porch object with the nearest horizontal edges in the background image. On frame 231, move the TrackerNull on its X and Y axes to match up the rear edge of the second porch step with the corresponding line in the background image. Cross-check this with the alignment of the front edge of the object's step and the estimated front edge of the step in the background image. If you can't get the two alignments to reconcile, you may need to change the PorchPivotNull scaling and try again.

From this point on, you will need to simultaneously tweak the PorchPivotNull scaling on all three axes, the TrackerNull X and Y position, and the Porch rotation on all three axes. An error on one variable affects the other seven. That may seem like a lot of variables to juggle at once, but remember that any remaining changes should be very small. By this time, any remaining errors in alignment, scaling or position should only be the width of a pixel or two on your screen. When you have frame 231 matched to your satisfaction, do the same for frame 265. Because frame 231 marks the end of the camera's movement in several axes, I recommend that you copy the settings for Porch rotation, PorchPivotNull scaling and TrackerNull position to frame 265. This will accomplish over 95% of the necessary tweaks for you; the remainder are relatively easy.

Now you need to step through the entire animation, looking for the frames where discrepancies peak. Wherever the match is off enough that it will show in the final rendering, create a set of keyframes to correct the problem. Some variables will require more keyframes than others. OriginNull and PorchPivotNull scaling should require the fewest (and identical number of) keyframes. Porch rotation will require slightly more keyframes. TrackerNull may require a key for every frame after frame 211, when the tracked motion file runs out. The XY plane movement of TrackerNull is the highest-frequency variable in this match move, and therefore the one that you will spend the most time and effort tweaking. Your manual keyframing between frames 212 and 265 should look very much like the tracker motion of frames 1 through 211, with tiny changes on most frames and the occasional larger jump.

When you have completed the move match, add an animated character or other moving LightWave element, parented to the Porch object. This doesn't have to be complicated, just something that casts a moving shadow over the steps and upper surface of the Porch object to show how well you've match moved the shadow-catching surfaces. If you like, use Load From Scene to add objects, motions and lights from the LampHop.lws scene file on the CD-ROM. This is a desklamp object that I animated to hop up the porch steps. After you Load From Scene, you will need to delete the default Light, leaving the SunLight imported from the LampHop scene. You will also need to parent both the desklamp object and the SunLight to the Porch object so they inherit the match move. You may want to leave Polygon Edges on to check the first set of renderings with the character added; after that, turn off Polygon Edges so you can accurately judge the lay of the character's shadow over the edges of the Porch object.

When you are satisfied with the test render results, render the entire image sequence. If you plan to composite a subtractive merge in Digital Fusion or other compositing software, you can leave the PorchSurface and background image as is. Your results should look something like Figure 6. If you prefer to comp an additive merge, you should drop the image maps from both PorchSurface and background, resulting in a completely black frame except for the desklamp.

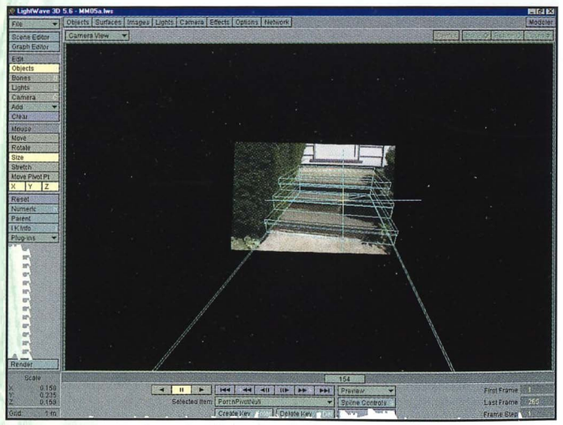

Figure 4: Porch aligned with horizontal background lines in frame 154.

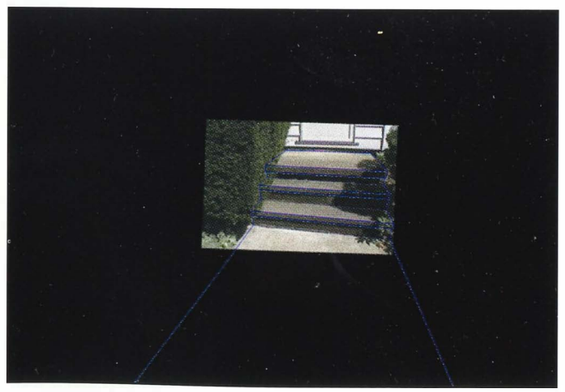

Figure 5: Rendered test of frame 154. Note how the blue polygon edges make evaluating the match much easier.

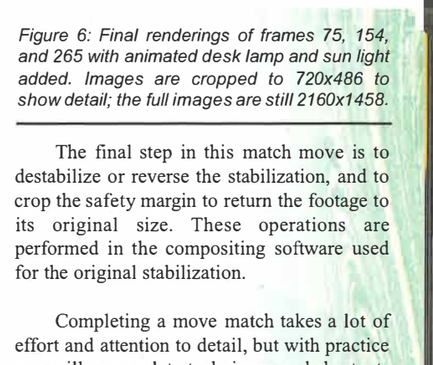

Figure 6: Final renderings of frames 75, 154, and 265 with animated desk lamp and sun light added. Images are cropped to 720x486 to show detail; the full images are still 2160x1458.

The final step in this match move is to destabilize or reverse the stabilization, and to crop the safety margin to return the footage to its original size. These operations are performed in the compositing software used for the original stabilization.

Completing a move match takes a lot of effort and attention to detail, but with practice you will accumulate techniques and shortcuts that make it faster and easier. Once you have a match move locked down, you can use it to render separate shadow, ambient, specular and Z-buffer passes. This can greatly increase your options for compositing effects, and can make comping LightWave characters seamlessly into the live footage much easier.

Doug Kelly is Editor of Keyframe and a freelance writer and animator. You can contact him via email at dakelly@megalink.net or through his website, www.megalink.net/~dakelly. This column is excerpted and condensed from his latest book, Digital Compositing In Depth, *available directly from Coriolis Press at www.coriolis.com and from major bookstores.*

Up Close and Personal

by Jason Goodman

When it first arrived on the scene at the dawn of the 1990's, the original Video Toaster was unlike anything anyone had ever seen. It was the granddaddy of all desktop video. For the first time, an inexpensive home computer could be transformed into a viable broadcast production tool. It put big broadcast equipment manufacturers out of business while putting countless independent producers into business. I would go so far as to say that it came dam close to being a household name. You may be surprised to know that a lot of those original Amiga Toasters are still in use in broadcast and industrial video production today. But 1989 was a long time ago. The Toaster/Flyer combo has been showing its age for quite a while. Developing for the Amiga is a technological and financial dead end. Despite this, NewTek has continued to support the Amiga Toaster and keep the dream alive. That was then.

Now we stand at the dawn of a new decade, and of course a new millennium. HDTV is on sale at Circuit City and $1500 iMacs have video editing and non-linear capabilities that rival the Toaster/Flyer. All of this commotion screams for NewTek to create a new generation of Video Toaster. Something like a Video Microwave. Something that would set the computer graphics industry on its ear once again and send the competition running back to the R&D lab.

Cross-platform Toaster compatibility has been a big issue since the Amiga days. Unfortunately for NewTek the Video Toaster's popularity surged just as the Amiga and Commodore's fortunes plunged. Early demands for the Toaster to expand beyond the realm of the Amiga led to the development and brief existence of a Toaster/Mac II interface. It allowed you to basically use an expensive Mac OS machine as a remote control for an Amiga (who came up with that idea?) I think it also allowed you to import Photoshop images via Ethernet, but it was hardly a Toaster for Mac users. Despite this early penchant toward the Mac, and the current success of LightWave 3D on the Mac, there have always been rumors only of a Video Toaster for Windows. Finally in the fall of 1999 NewTek has delivered the next generation Video Toaster. Of course this new development begs the question, can NewTek duplicate the excitement, originality and market-stirring price performance standards set by the original Video Toaster? That is what I have been investigating for the past several months.

The Video Toaster NT is highly dissimilar to its legacy namesake. The new Video Toaster is revolutionary in its design and its approach to processing digital video on a computer. In many ways it is ahead of its time. Instead of using fancy hardware codecs and proprietary on-board custom chips, it relies on smart design, high performance software and fast CPUs to provide D1 quality uncompressed capture, output, effects and transitions. However, that statement needs to be qualified with an "eventually." For now, many of the Toaster's features are not yet activated. At NAB NewTek showed off some very exciting and high quality DVE-type effects far superior to those of the original Toaster. They threw terms around the show floor like "unlimited video streams" and "real time effects" more freely than promotional T-shirts. While this all sounds impressive, these features are not currently available. Where the original Toaster was a DVE unit and switcher long before it got its non-linear editing chops, the new Toaster takes the opposite route. It is currently primarily a non-linear editing product with no real-time switching or real-time effects capabilities. In-sync's Speed Razor is the editing interface, and all effects must be rendered. NewTek also includes its video paint package, Aura, as well as a scaled-down version of LightWave. The Video Toaster NT is the most cost-effective system for getting D1 quality uncompressed video into and out of your computer. It has remarkable promise and potential, but for right now it is experiencing some teething pains. This product has some serious limitations that will make it very difficult (if not impossible) for many professionals to use.

I compared the features and capabilities of the Video Toaster to the Matrox DigiSuite LE. The Matrox product retails for $3995, a bit more expensive than the Toaster, but it offers an impressive array of features. Although the DigiSuite LE does not support uncompressed capture or playback, it still gives the Toaster strong competition in many areas. The DigiSuite LE comes bundled with Adobe Premiere RT, CG Inscriber and Boris FX. Curiously, the NT Toaster has no dedicated CG application. Character generation was one of the original Toaster's strongest suits. The Matrox board offers variable compression with a low, low 1.3:1 minimum. This is the same compression ratio used by broadcast formats like Digital Betacam. By providing two hardware codecs on the board, DigiSuite LE allows real-time effects and transitions.

One huge advantage the Matrox product offers is the ability to work in a variety of non-linear editing environments. I like Premiere, but DigiSuite LE is also compatible with Speed Razor, Discreet's edit*, IMC's Incite, and Ulead's Media On Line Express. This flexibility is great. I find one of the Toaster's most frustrating limitations to be its monogamous marriage to Speed Razor. NewTek is working on its own editing interface that is expected to take advantage of the Toaster's DVE capabilities, but it is unclear when that will materialize. So far there has been no official word on the Toaster's compatibility with other applications.

Another important advantage of the DigiSuite LE is its native file format. DigiSuite LE creates AVI files, but the Toaster uses a proprietary format called RTV. For now, the RTV format is only recognized by programs in the Video Toaster suite. This does not include After Effects, Premiere, Media Cleaner Pro, or any other video/media applications. Users who want to bring Toaster clips into these applications will need to first export an AVI from Speed Razor or Aura. This presented serious problems. As of software version 1.02, Speed Razor cannot output AVI files longer than one minute. Additionally,

Figure 1: Errant field insertion problem with VTNT.

AVI files coming out of Speed Razor could not be compressed by Media Cleaner Pro without serious amounts of screen garbage appearing in the clips. They had to first be processed through Premiere. This was also necessary with clips longer than one minute. Premiere can maintain image quality by outputting uncompressed files, and Media Cleaner can properly compress this output. However, this is a huge flaw that makes putting long form Video Toaster material on the web, CD-ROM or DVD very difficult. Creating high quality QuickTime footage from the Toaster takes three times as much space and time as it should.

When I tried to call in-sync to get technical support they told me to call NewTek. Frustratingly, when I called NewTek they told me to call in-sync! I quickly brought this catch-22 loophole to the attention of the higher-ups at NewTek, and I doubt that you would run into the same problem I did. NewTek is scrambling to get its tech support staff up to Toaster speed, but it has not been a smooth journey so far.

By using the more standard AVI format, clips captured and edited with DigiSuite LE work in many other applications (including After Effects and Digital Fusion) seamlessly. However, I did run into problems using DigiSuite LE's AVIs in Media Cleaner Pro. The AVI files created by DigiSuite do not interleave audio within the file. For every AVI there is a corresponding WAV (the Toaster does this too), and this can create problems for some applications. I believe the problems I saw in Media Cleaner are due to QuickTime. Media Cleaner uses QuickTime to read all AVI files. The QuickTime player also did not recognize the audio in files output from DigiSuite LE. As with the Toaster, DigiSuite LE's output had to be processed through Premiere before it could be compressed with Media Cleaner.

Another important element of the DigiSuite LE system is its software codec. Available from Matrox's web site, this piece of software allows any PC to use and edit media from the DigiSuite LE. PCs using Aura or LightWave 3D can recognize RTV clips, but setting this up is not as straight forward. Identifying the necessary software component required a call to NewTek tech support. Neither manufacturer offers a way to use these files with a Mac. While this is frustrating, it is not surprising. I feel that it is technically possible. If Matrox were to write a QuickTime codec (much in the way Intel has with its Indeo format) perhaps Mac users would find it more appealing. However, this seems unlikely at best. Currently the only video capture system that I am aware of that is fully cross-platform compatible is from Pinnacle Systems. The Targa 2000 line of video products is compatible with both Mac and PC. The more expensive RTX board offers similar features to the DigiSuite LE with dual codecs on board, real-time 2D effects and variable compression with a low minimum.

It is important for the prospective buyer to know that both NewTek and Matrox have "qualified" specific systems and components for use with their products. You can see these recommended systems on their respective websites. For the best results it is highly recommended that you stick to the qualified system specifications. I deviated from these specifications with both systems. This complicated troubleshooting and tech support. Specifically, neither manufacturer had qualified the drive subsystem I was using. My motherboard had been qualified by Matrox, but not NewTek. My audio card was also not approved by NewTek, but DigiSuite LE's audio is built in. This deviation resulted in no problems with the DigiSuite LE. It may have contributed to some of the problems I had with the Toaster relating to capture and playback but not AVI rendering.

The photo inset (Figure 1) shows a sample of the main problem I had with the Toaster when capturing footage. The captures were plagued by errant field insertions. In the example you can see that in the midst of a scene, a field from the previous shot is inserted (randomly) creating a highly distracting effect. This made capturing footage of any length a major challenge. NewTek has been making huge efforts to try to accommodate me and fix this problem, but so far none have worked. They even put me in direct contact with the top-level programmer on the Toaster team! Most recently they offered to have me send my machine to them for testing. I did not take them up on this offer as it was unclear when the machine would (or if it in fact could) be fixed. I do however commend them on these efforts.

For many users, uncompressed video may be a disadvantage in disguise. Uncompressed digital video offers the highest image quality and is best suited to applications like multilayer compositing or putting video on the Web, CD and DVD. The problem is it takes up so damn much space. NewTek's argument is that hard drives are cheap and you can just buy more. This is true, but in a studio environment, no computer is an island. Moving uncompressed video around a network can become a problem. Additionally, by having no options for compression, a user is forced to create videos that are maximum space hogs even for projects that may not require the ultimate video quality. By taking advantage of variable compression, DigiSuite LE gives the user more flexibility. DigiSuite LE also offers "uncompressed internal processing" for maximum quality. I am not sure what that means, but the board did produce excellent results. I have also included sample images of uncompressed Toaster footage and compressed DigiSuite LE footage for comparison. Unfortunately, the footage I used came out of an Avid and had been previously compressed. However, both the Toaster and DigiSuite LE had equal disadvantage in using the same footage. You may notice some artifacts in the 800% blowups, but I could see no difference in the uncompressed Toaster

Continued on page 28

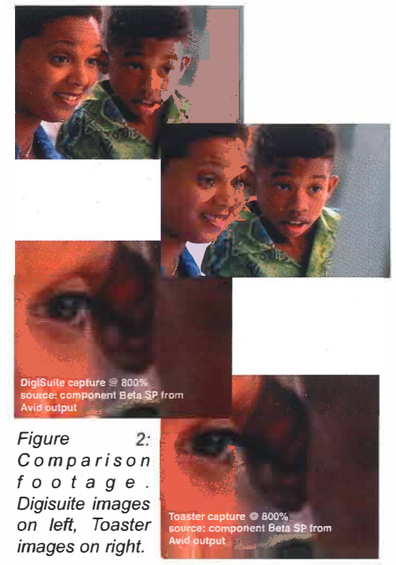

Figure 2: Comparison footage. Digisuite images on left, Toaster images on right.

64K Limit, Morph remodel, Skinning,

by David Warner

If you read my comments in the previous issue of *Keyframe*, you know that I was highly skeptical of NewTek's announced release date of "3rd Quarter 1999" for LightWave version 6.0. Well, at the time I'm writing this, we are one month past that deadline and there is still no new information on a release date for LightWave 6.0.

Many people on the Internet are complaining about this and giving NewTek a lot of grief about missing an announced ship date, but frankly, I'm relieved that LightWave 6.0 has not begun to ship yet. Sure I'm excited about getting my hands on all those cool new features, but the software I saw demo'd at SIGGRAPH '99 was simply not ready for public release. Even now, at the end of October, I'm guessing that NewTek has a lot of work left to do. I hope that by the time you read these comments we will have a much clearer idea of when LightWave 6.0 will be available, or even better, we'll all be too busy playing with it to even care about missed ship dates. However, it's important for people to be patient and accept the fact that when the new version is ready, it will ship. NewTek would not be doing any of us a favor if they released buggy or incomplete software, so let them take their time and get it right.

Another highly anticipated LightWave tool is on schedule and has begun to ship when everyone expected it to. Project:messiah, the amazing new character animation plug-in from Station X Studios began shipping in late October and this is news that could very well be more exciting than the release of LightWave 6.0! Messiah's feature list has grown even longer since its unveiling at SIGGRAPH. A tool that I considered to be a 'must have' several months ago has actually grown more impressive. Station X has a real winner on their hands, and I'm betting that even with a $995 price-tag, this plug-in is going to become extremely popular among the LightWave community. I can't wait to see what these guys come up with next!

Now here is this issue's selection of questions and answers from the LightWave mailing list.

Q I have created a highly detailed object in Modeler and when I try to save it, I get an error message that says "Too many points to save in file - File cannot hold more than 65,536 points." Why am I getting this error and how can I still save my object?

A The 65,536 point object limit is an annoying problem that has plagued LightWave for a long time, but I hope that by the time you're reading this it will be a thing of the past. Basically, the 64k point limit is a result of the way LightWave handles mathematical calculations, which is based on 16bit precision. 16bits of data equals a maximum number of 65,536. In more practical terms, it's a limitation in LightWave's object file format that can be a real pain.

This limitation will be eliminated with the release of LightWave version 6.0, but until then there are a number of ways to get around this problem. One way is to simply use Modeler's Cut & Paste functions and the multiple work layers to cut your object into pieces. Then save each individual piece as a separate object and rejoin the pieces in Layout by parenting them all to a null object. The individual pieces will maintain their positions relative to each other. If you cut away the individual pieces based on different surface names, there won't be any noticeable seams where the different objects connect.

Perhaps an even better way to get around this problem is to use a point reduction plug-in. This tool will analyze the object to find areas of low detail where extra points can be removed without affecting the appearance of the object. There are point and polygon reduction plug-ins included with version 5.6 of Modeler, but a much better plug-in is available for free on the Internet. QemLoss2 does a fantastic job of reducing an object's point count and best of all, it's FREE! You can find a link to this and many other free LightWave plug-ins by visiting one of the best LightWave websites on the net: www.flay.com.

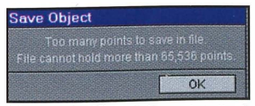

Figure 1: The 64k point limit error message.

Q How can I prevent color banding when using the Steamer plug-in?

A One way is to raise the number of Samples that Steamer uses when generating the effect. This is a value that you can change in the Steamer plug-in's interface. The only problem with raising the number of samples is that it can significantly increase the render time for your image. If possible, a better solution is to pull the camera away from the Steamer effect, because color banding in a Steamer effect usually only occurs when the camera is too close.

Q I like using Area Lights to give my renderings a certain look, but the render times can be a bit too long for production work and sometimes the final render can look very grainy. Is there anything I can do to make the image quality look better and speed up the render times when using Area Lights?

A You can eliminate the graininess from the image by increasing the anti-aliasing settings in LightWave's Camera panel. A setting of at least Enhanced Medium will make your image look a lot better, but the render times will increase significantly. If you want to eliminate graininess and decrease render times, then you need to decrease the size of the area light so the fuzziness of the shadow is not as soft. The softer the edges of the shadow, the higher the probability of graininess in your shadows.

If you're familiar with the concept of multi-pass rendering, then you might consider rendering your area light shadow separately from the rest of the scene, then blur that layer when you composite the pieces of your final image together. Another alternative is to not use area lights at all, but try to fake their effect by using strategically placed point lights and the Soft Filter option in the Camera panel. And last but not least, owners of Gaffer (from Worley Labs) should check out the powerful area light controls available in that commercial LightWave plug-in.

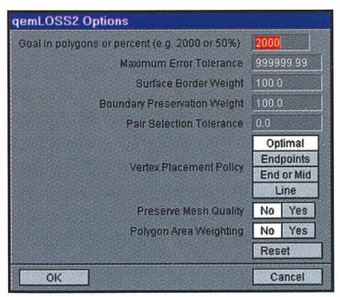

Figure 2: The QemLoss2 interface.

Q I have set up a morph between two objects and it is working correctly, but now I need to add a small bit of detail to the object. I don't want to have to recreate the morph target object all over again, so I'm wondering if it is possible to add geometry to the morph source and target objects and still have them morph correctly?

A In theory, if you do the exact same thing to both objects, they should still morph correctly. However, the general rule of thumb is that you can morph points around as much as you want without messing up the morph, but if you add any geometry then morphing will not work anymore. If you do add geometry to each of the objects and the morph does break, there is a free plug-in available on the internet called "LWO-Conv" that will analyze morph source and target objects and reorder the points of the target so that the morph works correctly. This is another free plug-in that can be found by visiting www.flay.com

It should be pointed out that the endomorph feature in LightWave version 6.0 will eliminate these kinds of morphing problems, so you can set up a morph between two or more objects and then add new geometry to the morph source object without breaking the morph. This will be a HUGE improvement in how morphing is handled in LightWave!

Q I obtained various cross-section images of a complex real-world object and I thought that I could use these images to model cross-section splines in Modeler and then Skin them together to create the whole object, but when I apply the Skin function, the object's mesh is deformed and does not look the way it should. The Skin function appears to twist the splines somehow, so that the new polygons are non-planar. Why is this happening?

A The main thing to watch for when using the Skin function is to make sure that the spline curves you are skinning have the same number of knots and that the knots line up in a correct order. If you try to skin splines together that have a different number of points or a different point order, then you get a lot of errors and other undesirable results.

When creating spline curve cross-sections for your object, it is usually a good idea to create one spline that represents the most detailed part of your object's cross-section, making sure that it has an adequate number of points (knots) to represent that part of the cross-section. Copy & paste that spline curve as many times as you need to represent the rest of the object's cross-sections. Use the Drag tool to move each spline curve's points around, repositioning them as needed to represent each of the other cross-sections. Once the curves are reshaped you can apply the Skin function. You should get much better results this time.

Q How can I save the shape of an object that is being deformed by a displacement map?

A In the Objects panel, click on the Save Transformed button. This will save a copy of your object with the modified point position data. It works for anything that changes the shape of an object: a displacement map, bones, morphing, a displacement plug-in, or even a new keyframed location. It does not, however, save the shape of an object created with a Clip map. Also keep in mind that if you have any textures applied to an object, using the Save Transformed function may have some nasty effects on the appearance of your object. The texture size/position data remains unchanged, but the shape/position of the object itself is changed, so texture maps may not be positioned correctly on the transformed object.

Q I've noticed that if I'm rendering a long animation and I leave the computer alone so it can do the work, after about a half hour, the render times get really big but very little has changed in the animation, at least nothing that would make LightWave take so much longer to render. While doing some tests, I found that LightWave actually rendered the animation much faster if I kept using the computer, instead of leaving it alone. What could be causing this?

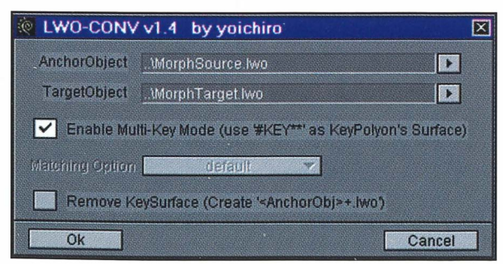

Figure 3: The LWO-CONV plug-in interface.

A Many new computers have a Standby mode for many of the components inside your computer, including the hard drive, monitor, and even the CPU itself. This Standby mode tells the computer to put these different components to 'sleep' if the computer is not used for a while. This saves energy and wear-and-tear on your computer, but it can also cause big problems for a program like LightWave. If the Standby mode for your CPU kicks in while you are trying to render an animation, then you are losing a significant amount of horsepower and your animation will take much longer to render than it should.

Windows users can find out if Standby mode is turned on and disable it by going into the computer's BIOS, which is the initial bit of programming that tells the computer what to do when you first turn it on. Different computers have different types of BIOS versions so these settings are accessed and changed differently for each computer. You need to restart your computer and watch the initial couple of boot screens very closely to figure out how to access your BIOS. Usually you press a key on the keyboard, like F1 or ESCape or DELete, during the first stages of startup. This will open up a set of menus that allow you to access different settings in your system's BIOS. You want to look for a Power Management menu, which is usually where the Standby mode is turned on/off. Be very careful about changing other settings in your BIOS, because it is possible to change the settings in such a way that your computer will no longer work properly.

Macintosh users can also disable the various Standby functions by going into their Power Management control panel, but this is a problem that does not usually affect the Mac community. Once you disable the Standby mode for your CPU, you can leave your computer unattended and still have good render times.

That does it for this installment of On Topic. For more LightWave tips and tutorials, be sure to check out the new Tutorials section on my website at http://www.tv3d.com.

For more info about David Warner or TutorVision 3-D CD, his series of multi-media tutorial CD-ROMs for LightWave. The newest release, Volume 3, contains more than 6 and a half HOURS of video and a total of 26 different lessons for LightWave! For more info, please visit http://www.tv3d.com or send some e-mail directly to Dave at dwarner@tv3d.com

Drawing Bitmaps in LWPanels

by Brenden Mecleary

A 'cool' user interface (UI) for a LightWave plug-in often makes the difference between the plug-in being taken seriously and it being considered a hack. For freebie plug-ins, this is really not much of an issue. People get what they pay for, and if they don't like the UI on a free plug-in, they just have to live with it. However, commercial or shareware plug-ins must make a good first impression. A 'cool' UI - and I use that word deliberately, instead of 'good' - often has a strong influence on a user's purchase decision. This is doubly important for LightWave plug-ins, since we are selling our software to artists. This is often a painful truth to face, since writing the UI can be exceedingly dry, tedious work, and a good UI for a complex LightWave plug-in can easily comprise 50% of the plug-in's code.

In this column, I'll briefly explain how to use bitmaps to add some graphical pizzazz to your plug-in. A really nice way to establish a consistent look for your plug-ins is to brand them with a logo. The same techniques can also be used to create custom controls, show preview information, or illustrate your UI more effectively.

Unfortunately, LWPanels has no native support for bitmaps. This means you will have to render your bitmaps within your own code. There are several ways to go about this, but by far the most common way is to draw the bitmap pixel by pixel using LWPanels's drawing functions. Keep in mind that this approach is very slow. If you require real-time updates, you should look elsewhere. I suggest an OpenGL control. The price you pay is that it's not quite as straightforward (if you don't know OpenGL, it's up to you to learn), and it's not quite as reliable - at least, not for LightWave.

There are two possible places to draw bitmaps. Panels and controls both support user drawing routines, so you can use whichever you like. Controls are typically nicer for drawing simple bitmaps and preview windows. It's easier to trap input for controls, and they also give you support for different states: enabled, ghosted, and so on. If you need to perform broader bitmap operations - say, something that draws around a few controls at once - you might want to use a panel draw callback.

The controls of choice for drawing bitmaps are the area control and canvas control. They have minor differences: canvas controls do not receive event notifications (area controls do), and canvas controls have a slightly darker background.

Once you have created the control you want to house your bitmap, you must assign it a draw callback function. Control draw callback functions must have this prototype:

```
void ControlDraw(LWControlID idCtl, void * pvUserData,
DrMode drMode);
```

LWPanel.h defines this to be a "CtlDrawHook", then never makes another reference to the type ever again. The first parameter is an LWControlID, a type that, if you look through LWPanel.h, would appear completely useless, since there is apparently no way to get anything useful from it. Although it's not obvious from the documentation, an LWControlID is actually an LWControl pointer. So, your callback should really look like this:

```
void ControlDraw(LWControl * pCtl, void * pvUserData,
DrMode drMode);
```

Alternatively, you can use a cast to convert the LWControlID to an LWControl pointer when you get inside the function. In this case, it's such a silly waste of time doing it the 'right' way (that is, matching the prototype in LWPanel.h) that I prefer to just change the function signature.

To attach your callback to the control, you must set its CTL_USERDRAW property. Although LWPanel.h provides macros for setting the user event handler, LWPanel.h does not do anything to help you set the draw function. Fortunately, it's easy to define an inline function to do this:

```
inline void LWCtlSetDraw(LWControl * pCtl,
CtlDrawHook pfnDraw)
{
    LWValue lwvInt = {LWT_INTEGER};
    ival.intv.value = (int)(pfnDraw);
    pCtl->set(pCtl, CTL_USERDRAW, &lwvInt);

    return;
}
```

As soon as you create your bitmap control on the panel, call this function to assign it your custom draw function. It's important to note that your custom draw function is called in addition to the default drawing behavior of LWPanels. For example, if you add a custom draw function to a button, you do not have to rewrite the code that draws the button. You just get the opportunity to scribble all over that button every time it's repainted.

Your custom draw function is where all the real work goes on. You have to increment through each pixel in your bitmap and draw it to the control. I recommend that your control and your bitmap be the same size just to keep this process as uncomplicated as possible. Actually drawing the pixels is rather easy, and can be accomplished with a pair of loops like the following:

```
for(unsigned int uRow(0); uRow < uBitmapH; ++uRow)
{
    for(unsigned int uCol(0); uCol < uBitmapW;
++uCol)
    {
        //Draw pixel here.
    }
}
```

You will need several things before you draw:
1) The LWPanelID of the panel containing the control.
2) The X / Y location of the control.
3) The DrawFuncs structure, which contains (surprise!) functions for drawing.

All of these are available directly or indirectly through macros in LWPanel.h. CON_PAN will return the LWPanelID for the given control. The control X and Y positions are available through the CON_HOTX and CON_HOTY macros. You should not use the CON_X and CON_Y macros, because they return the X / Y position of the entire control, including the label. The CON_HOTX and CON_HOTY macros return the X and Y positions where you should start drawing. To get a DrawFuncs structure, use the CON_PANFUN macro to get an LWPanelFuncs pointer for the panel. A pointer to a DrawFuncs structure is available inside the LWPanelFuncs structure.

When you are ready to draw pixels, you need some pixels to draw. This is where you must do some footwork. Although most platforms provide some method to embed binary resources in an executable, none of these methods are portable. Unless you feel like writing a platform-independent wrapper of some sort for that service, you will have to store your bitmap as an initialized C array with some sort of converter. There are a couple of converters already available, so it would be a smart choice to use one of those. Bob Hood has written one that works with TGA images, available on www.lightwave-outpost.com. Also, I have included a converter on the Keyframe CD-ROM that uses the Windows Bitmap format. If you still intend to write your own converter, consider using the RAW format that is available in newer versions of Photoshop. This format is what it sounds like - raw data. There is no header of any sort, just the bitmap. Writing a converter from RAW to C would be very easy. If I weren't so cheap as to still be using Photoshop 2.5 LE, I'd have used that format myself.

Once you have your bitmap in a format the compiler will accept, write your draw function. If you've been thinking ahead, you may have reached the conclusion that since one bitmap is like the next, you can isolate the bitmap drawing code and re-use it for multiple controls. That's exactly what you should do. I use a simple structure like this:

```
struct LWPanelBitmap
{
    unsigned int   m_uWidth;
    unsigned int   m_uHeight;

    typedef const unsigned char RGBPixel[ 3] ;
    RGBPixel                * m_paBitmap;
} ;
```

Create one of these, set it as the control's userdata pointer, then set the control's user draw function to a generic bitmap draw function. If you need a different user data structure, you can write a custom draw function to call the generic bitmap draw function with the appropriate parameters. You could also make the first field in your own structure be the LWPanelBitmap you want, and the generic draw function will be none the wiser. That route is a little riskier, however. You take fewer chances by using a tiny custom draw function that calls the generic one. Static class functions come in handy for this sort of thing.

Your generic bitmap draw function should be relatively easy to implement. Go through every pixel in the bitmap, and draw it to the panel. Be sure to offset your X and Y locations by the X and Y position of the control, since the panel draw functions use a coordinate system relative to the panel, not the control.

It is also possible to implement a primitive transparency system while drawing the bitmap by adding a 'transparent color' field to the LWPanelBitmap structure. You will need to revise your draw function so it won't draw any pixel having the 'transparent' color. To get really fancy, you can add a second bitmap that gets drawn when the control is disabled. By checking the DrMode argument to your draw function, you can decide which version of the bitmap to draw.

Bitmaps are a great way to spice up your user interfaces without too much effort. The branding effect alone is easily worth the time invested. The same techniques also come in handy as you add more sophisticated UI elements to your plug-ins' interfaces.

Brenden Mecleary writes funky Windows apps at work, and polishes up his game of Soul Calibur in his spare time. The people at Electronics Boutique are starting to get suspicious about him spending so much time there. Comments, questions, and anything else, e-mail bman2@ix.netcom.com.

by Charlie Ramos

The Metamorphosis Part I credits:
Directed by: Charlie Ramos
Produced by: Filmtel, Charlie Ramos
Screenplay: Charlie Ramos, Deirdre Scully
Voices: Karl-Heinz Teuber, Deirdre Green, Brian Douglas, Deirdre Scully
Music: Jaume Carreras
Camera & Lighting: Begonia Colomar
Animators: Feba, Charlie Ramos
Sound Design: Marcel Vilar
Sound Mixing: Emili Lleonart, Marcel Vilar
Color: Fotofilm S.A.E.

"This film is a segment of Kafka's novel beginning when Gregor wakes up to find himself changed into a giant bug, and ending when his family and boss convince him to open the door and thus discover his change. The reason I chose to do this segment (aside from money and time constraints) was that I felt this was the strongest segment, both visually and story-wise.

There are four parts altogether in the original version, beginning with a dream sequence that visually tells what Kafka explains narratively. In this sequence I wanted to show the sense of anxiety Gregor feels as his family surrounds him and tries to find out why he won't get out of bed. I also wanted to establish their relationship to him through their individual reactions to the situation. Like most of Kafka's books, the story does not have a happy ending. Gregor, after being abandoned and becoming a source of shame for his family, finally dies of sadness, disease, and neglect.

I guess my name is everywhere in the credits, but I think that most short computer film projects tend to be that way. I like all aspects of this media. I thought about composing the music, too, but luckily found someone willing to do it. I think that the fact that it is all in your computer makes you feel like you could possibly do everything, although this is not a good way to approach a film project. Although someone needs to have an overall vision of what the thing is about, the more feedback you get from other contributors the better.

I didn't have a specific intent when I started this short, so I suppose it was just for fun, the same feeling as planning a clubhouse when I was a kid. The best part was just dreaming up all the impossible things it would have, such as Star-Trek-style sliding doors. *The Metamorphosis* was just a dream. My wife had a Super 8 camera, so we were going to print out the images in black and white and film them one by one. I didn't really consider this as a serious project for festivals until I started working at Filmtel, and they showed an interest in helping me produce it. It was like a gift from heaven. Filmtel has pretty much everything as far as equipment; you can walk in with an idea and leave with a film.

If you count pre-production (which I did in my spare time), this film took 3 years from start to finish. Actual full-time production and postproduction took about seven months. I was given the OK to make the short in December of '98. The main reason for the decision was that Filmtel's 3D animation department was still new and they had nothing for me to do yet. I hadn't considered offering Metamorphosis at first because it seemed a little dark for what they had in mind, but eventually I realized that it would be stupid not to take the opportunity. I set up a small studio at home, convinced my wife, Begonia Colomar, to do the lighting and camera work, and hired a stop-motion animator, Feba, to help me with the animation. Neither of them knew the software, but each had enough

experience in their field to catch on immediately. A month later, we were sending CDs full of images for color tests. Besides studying film, Begonia had worked at a film lab and was able to translate what I wanted through that very important last step.

When I started the initial modeling, I was using a Pentium Pro-200 with 125MB RAM and one copy of LightWave 5.0. Later on, I purchased 3 more copies of LightWave 5.0, one upgrade to 5.5, and a Pentium II-400 with two processors, 125 MB RAM each, and a 9GB hard drive. Later I purchased a final Pentium II-400 with 1 processor and 125MB RAM. These were all the computers used for the bulk of the production.

The characters were animated using Puppet Master. We divided the work into body and facial animation. The first step was using Magpie. I would dope out the dialogue and write notes about the animation. I would then talk it over with Feba and he would draw out the main poses on paper. After looking at them, we worked out any problems they might have. Next Feba was given the scene file, which included the characters and a very basic version of the background , along with a rough camera angle. Once he had done some of the keys, he would pass it on to Begonia. She would do the final camera angle, and then he would finish the shot. During this time I mainly concentrated on the characters' facial animation and shot continuity. After all animation was complete, Begonia would light the background and characters. The textures where done using Photoshop 5.0 and extra surfacing and lighting effects were done using Gaffer and Bloom. Rendering was done at night using Screamernet. All in all, the process went very smoothly.

Editing was done using Premiere 3.0. After the shots were composited and rendered, we copied them to CDs and sent them to Filmtel. All the postproduction including titles, color correcting and final touch-ups were done using Kodak's Cineon software. This step was actually one of the most important, because the colors that you see on your monitor have nothing to do with the colors that will be projected in the theater. It was only a matter of passing the shots through a couple of parameters; after two or three worked, we didn't even have to look at the rest. The images were then transferred to film using a Solitaire 5 and sent to the lab for processing. Once the color timing was done at the lab, it was telecined back to video. Two versions had to be done at this time, one at 24 fps so the sound effects dialogue and music could be synced for film, and another at 25 fps PAL to send to festivals and friends. During this process we must have made a dozen trips to the lab to see how the final color timing was coming. Once we were happy with the results, the sound was added to the film and the final copies were sent out.

Metamorphosis has been shown at a number of film festivals, and a DVD version is being made. To date it has been shown at the SITGES Film Festival (Catalonia), Ft. Lauderdale Film Festival (USA, www.ftlaudfilmfest.com), Cinanima Animation Festival (Portugal, www.cinanima.pt) and Bilbao International Film Festival of Shorts and Documentaries. We also won the infografia en españa division of Art Futura (www.artfutura.org) in Seville, which is a kind of SIGGRAPH here."

Charlie Ramos can be contacted via email at charlieramos@hotmail.com, or phone at (34) 93 459 31 27. For more information about Filmtel, email info@filmtel.com, phone 93 494 74 00, or send mail to Filmtel, Edifici L'illa, Avda. Diagonal 549, 6a planta, 08029 Barcelona.

Las Vegas - February 16-24
Photorealism Series

Dallas - March 9-20
Character Animation Series

Chicago - March 29 - April 6
Photorealism Series

Topeka - Budget Workshops May 16 - June 2
Our entire curriculum in 18 intensive days at bargain prices!

Philadelphia - June 5-16
Character Animation Series

Phoenix - July 10-21
Character Animation Series

Columbus, Ohio - August 2-10
Photorealism Series

Topeka - Budget Week September 8-17
Character Animation Series

Salt Lake City - October 11-22
Character Animation Series

Atlanta - November 8-16
Photorealism Series

Photorealism Series
- LightWave 3D Essentials (3 days)
- Lighting/Camera Techniques (2 days)
- Surfacing Techniques (2 days)
- Compositing Techniques (2 days)

Character Animation Series
- LightWave 3D Essentials (3 days)
- Advanced LW 3D Techniques (3 days)
- Organic Modeling & Character Animation Essentials (2 days)
- Intermediate Organic Modeling & Character Animation (2 days)
- Advanced Character Animation: New L[6] Tools (2 days)

Washburn University
Division of Continuing Education
1700 SW College
Topeka, KS, USA 66621
lightwave@washburn.edu
785-231-1010, ext.1399
FAX: 785-231-1028

Cover Study - Photorealistic Leather Texture Mapping

In this tutorial I will describe how to create the surface groups, settings and texture maps to produce a realistically worn and weathered leather jacket. Many different factors contribute to making the material believable. The first factor to consider is the LightWave 5.6 mapping system. The likely candidates among the standard mapping procedures would be planar and cylindrical, because nothing in this project really lends itself to spherical or cubic mapping. Any planar mapping can use screen grabs from Modeler as painting templates in Photoshop. However, the cylindrical mapping will require Ernie Wright's Unwrap modeler plug-in, which you can download from http://members.home.net/erniew/getstuff.html . This plug-in unwraps and flattens a 3D mesh from Modeler to create a 2D black-and-white image file. Because the unwrapped image perfectly matches LightWave's cylindrical mapping coordinates, it can be used as a painting template for accurate texture placement.

Keeping the mapping options in mind, I split the mesh into its surface groups as shown in Figure 1, ending up with seven unique surfaces.

Figure 2 shows the first surface which I called GenkoJacket.srf. This surface will be cylindrically mapped and will need some detailed painting, so I need to know where exactly I am stroking the brush. This is the first instance to use the Unwrap plug-in. There are a few rules to follow for a successful unwrap. First, copy the GenkoJacket.srf from the main mesh and paste it to another layer. The newly pasted object usually needs centering; even if the mesh it came from was perfectly centered, the new pasted object more than likely has a new bounding area and thus needs re-centering.

With the object centered, I ran the Unwrap plug-in which opens up to show a dialog box, Figure 3. Knowing that the object is oriented in +Y, I checked the appropriate button and entered the desired Height and Width resolution, measured in pixels.

Figure 4 shows the results of the unwrap plug-in procedure. What the plug-in produces is a flat representation of the surface that lies on the mesh, showing the exact location of specific areas that are going to need detail painted into them. When the texture has been created using this template, it is an easy process to apply the map to model. By choosing cylindrical mapping and automatic sizing with a default wrap of 1.0, the texture should map perfectly to the mesh.

Figure 1: Jacket mesh split into seven surfaces.

Figure 2: Genko-Jacket surface.

Figure 3: Unwrap plug-in dialog.

The next few steps will show the five different attribute images that I painted, pointing out how each one is used in the resulting material. Every other surface which needs the leather material will need to have the same set of images painted: color, diffuse, specular, bump and sometimes a second bump image for stitch detail.

Figure 5 shows the first image to be created, which has a major importance to the overall look. If the color palette is wrong the surface is never going to look right. I found this darker brown tone easier to achieve than a lighter shade like tan. What I used was a piece of leather I had scanned and placed this over the wire template. At this stage it looks nothing like the jacket, it is obviously a flat scan. A critical requirement for the scan at this stage is that the leather's ridged grain has to appear to be the right size; anything too big looks false.

Once I had a scan that looked right, I enhanced the texture by taking a dodge brush and lightening the areas around the edge, making the seams appear to be worn through more flexing of the material. From the wire template I could see where the pockets and strap hoops are positioned. I changed the saturation around these areas to reinforce the worn look of the leather. Other color change details are a stitched ridge pattern laid on top of the scan to give the appearance of detail. Where there are light areas there are usually corresponding dark areas. I added these by applying a burn brush under the arm hole and in general those areas that would catch the most dirt.

Figure 6 details the painted bump map. I started with a plain medium gray canvas to give an equal base for positive and negative values (light and dark areas). Figure 2 shows that although some organic flow has been modeled into the mesh it still looks too evenly rigid. First I wanted to break up the base of the jacket by painting dark strokes that form the shape of a fold ridge. Once this looked correct, I painted the light areas, which to work right need some mid gray in between the two. What I was looking for here was to consistently paint graduated streaks. Painting light and dark on separate layers helps with this process, as it enables the artist to softly delete certain areas of each layer until each stroke has a gradient ramp between light and dark.

At this stage I applied the map to the mesh just to check that I was on the right track with my painting. The same mapping procedure used to apply the color map should apply here. It's easy to copy the values from the color to the bump texture panels. In the Color Texture panel, press Ctrl C to copy the settings; switch to the Bump Map panel, then press Ctrl V to paste the coordinates and other values. The bump map defaults to a 50% amplitude. I swapped the bump image for the color image, increased the bump strength to above 500%, left pixel blending on and changed the texture antialiasing strength to 0.25.

A render at this stage should show some fluidity to the base of the jacket. If the painted strokes don't look right, try varying bump strengths and tabbing between the painted layers to make adjustments where needed. When the base looked right, I applied the same techniques to the other areas of the jacket, especially near seams and around the pockets.

For applying the first bump map, extraordinary bump strength is used. However, there are still other bump details that needs to be applied to the mesh, namely the stitching pattern. As figure 7 shows, this is quite simple: a black canvas with line detail around the pockets and straps and general tailoring pattern that matches what was painted into the color images. The lines are painted as a mid gray on top of the black with a white broken line above the gray to give maximum scale through the bump. I applied this second bump map to the mesh at its own bump strength, which is nowhere near the extremity of the first bump map.

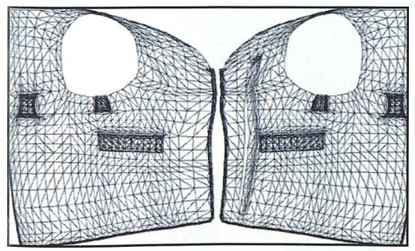

Figure 4: Unwrapped template of GenkoJacket surface.

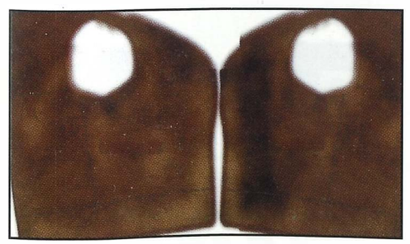

Figure 5: Jacket color map.

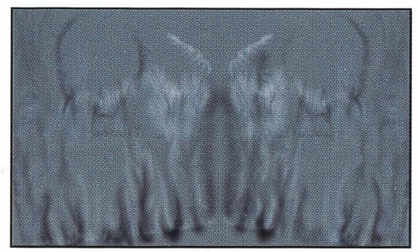

Figure 6: Bump map.

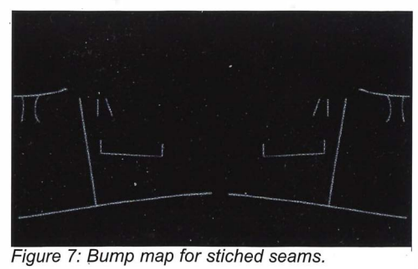

Figure 7: Bump map for stiched seams.

Painting the diffuse map was quite simple because I had the previous images for references. The bump image from figure 6, inverted, is a good basis for the diffuse map. I adjusted the contrast to even out the range of the tones. I applied it to the mesh and rendered it to gauge the effect. Generally, an even gray over the majority of the darker diffuse areas is needed. This allows other details to be painted on top and gives scope to add darker details under the arm and where the stitching would be. Note that the stitching is still the same gray range as in the original and not inverted. We want the stitching itself to be light and dark surrounding the stitch to give the appearance of dirt gathering. To achieve the light stitch area, I took the map shown in figure 7 and overlaid it, which showed the correct areas to paint the gathering dirt.

The last image in the set is the specular map, which plays a very important part in the overall look. This map should take into account highlighted details from all previous maps, including any stitch details and surrounding details at the pockets. I created the image by taking the color image, desaturating it and lowering the overall brightness to achieve a lower spectrum for the specular information. I then overlay the bump map to get lower specular values around the negative areas of the bump. This should result in a well-broken-up specular appearance to the texture which leaves the specular glossiness values which will appear around the affected areas. The glossiness value is controlled using Worley Labs' Gaffer plug-in, looking for a low sheen. Figure 10 shows the results of each image in turn.

Having achieved the right look for the jacket's largest leather surface, I applied the same techniques to the other leather areas on the jacket. For the remaining four leather surfaces I only needed the Unwrap plug-in one more time, for the belt which flattens in a similar fashion to the jacket body. All others can be created using screen grabs from Modeler in the relevant view to generate the templates for painting. When working on these remaining images, it is important to pay attention to where the stitching lines up. This aids in blending the surfaces together, giving the appearance of one flowing pattern over certain surfaces. In this case it is most apparent where the front flap joins the jacket body. Figure 11 shows each surface rendered from its mapping axis point of view.

With all the leather surfaces in place, three unique surfaces, the jacket buttons, the waist belt buckle and the inner jacket surface remain to be textured. The buttons are quite simple; choosing LightWave's marble procedure gives a good result. The desirable results are evenly distributed veins across all the buttons and a medium gloss specular highlight as seen in Figure 12.

The waist belt buckle carries a metallic appearance so a low diffuse with reflection options set gives a good base for creating this material. Gaffer can produce the correct specular attributes with linear anisotropy set to about 50%, as visible in figure 13.

This leaves the inside of the jacket. For the purposes of this tutorial I created a unique surface that covered the entire area within the jacket. However, for working practices when this outfit actually fits a character (as shown in Figure 15) there is no need to cover inside totally, just the visible areas like the base of the jacket. To get the geometry for the surface I selected the jacket body, smooth shifted by -3mm and then flipped the polygons. I applied a simple tileable silk texture (Figure 14) cylindrically mapped and used Gaffer to create a soft sheen across the surface.

Figure 15 shows the completely textured jacket on the Genko character for which it was tailored.

Graham McKenna is a 3D artist in Scotland. You can email him at graham@grahamprismatic.free-online.co.uk or visit his web site at http://www.grahamprismatic.free-online.co.uk for similar imagery and techniques.

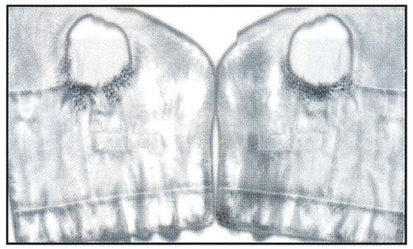

Figure 8: Diffuse map.

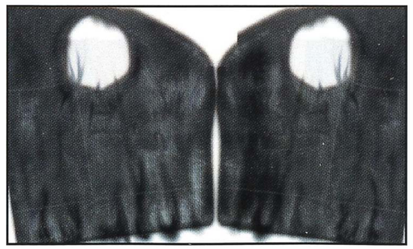

Figure 9: Specular map.

Figure 12: Button texture using LightWave marble procedural shader.

Figure 13: Belt buckle textured using Gaffer.

Figure 10: (left to right) color map; color and bump maps; color, bump and diffuse maps; color, bump, diffuse and specular maps.

Figure 11: Jacket front flap, jacket shoulder rims, jacket collar and jacket waist belt.

Figure 14: Tiled silk image for jacket lining.

Figure 15: Genko wearing finished jacket.

THE JOY OF SIX!
LightWave [6] – Experience the seduction of **LightWave [6]** ! Completely redesigned and mastered, with the artist in mind, for creating exceptional realism faster and easier than ever before! Powerful new tools include enhanced interactivity and scene editor, Skelegons, MetaNURBS in layout with point weighting, caustics, UV mapping, surface editing in modeler and so much more! LightWave 3D is the most used animation system on the planet. What are **you** using? **2495.00**

LightWave 3D Competitive Upgrades ..... **Call**
LightWave 3D Cross Platform Upgrades ... **Call**
LightWave 3D Upgrades ..... **Call**
Aura ..... **Call**
Aura Complete **699.99**
Inspire 3D with LD_Tools **429.99**
Graphics Studio ..... **Call**
Video Toaster NT **2995.00**

NewTek

Visit:
www.videotoasternt.com
Complete reliable, quality turnkey Toaster workstations are available.

LifeForms 3.5 Creativity in Motion

Life Forms 3.5 LE The latest version of the defacto standard for creating realistic figure movement for 3D animation. Its intuitive interface and powerful features will help you save time and be more productive than ever. Our new entry-level package gives you all the features of Studio without the application support of 3D content. Great for creating Quicktime movies, VRML output for the Web, input for Strata Studio or trueSpace **179.00**

Life Forms Studio 3.5 – All the great features of LE, plus support for many of today's top 3D animation packages. The Studio version includes PowerMoves 1 & 2, a collection of over 750 motion files (valued at up to $5,000) to get you moving fast! **449.00**

CREDO The Character Motion Company
www.charactermotion.com

LIGHTWAVE PLUG-IN/ LIGHTWAVE RELATED SOFTWARE

Alpha Channel Landscape Textures....*NEW!* ..... Call
Alpha Magic 199.00
Amazing Sci-Fi Textures 129.99
Animation Master 99 ..... 176.99
Animation Master/Simbiont Bundle 219.00
Aqua Pack 156.99
AuraFX Plug-ins 119.99
Bryce 4 ..... 185.99
Complete Effects Library 399.00
Creature Creator for LightWave ... 225.00
DarkTree Textures 289.00
DirtyUp 229.99
Dosch Textures Vol. 1 ..... 77.99
Enlightment (electronic version) .. 79.00
Hollywood FX ..... Call
Light ROM Super Bundle w/Leny ... 99.99
Lip Service 299.00
LW 3D*ROM 49.00
..... 289.00
..... 699.00
..... 489.00
..... 129.99
Puppet Master 299.00
Relativity 189.99
Seamless Textures 4 – Classic Stone .....
Seamless Textures 5 – Downtown Surfaces....*NEW!*.
Space Wars Digital Effects Kit ..
Tree Druid, Intel/Mac
Visual Effects CD Set
World Construction Set 4.5 Intel/Mac

dpsReality
The world's first Studio Digital Disk Recorder! Features compressed or uncompressed recording, alpha channel, balanced/unbalanced stereo audio I/O and DPS Virtual Tape File System for group rendering. The package contains the dpsReality player application, Digital Fusion DFX compositing software, and an audio/video breakout cable. **2249.00**

dpsReality Studio Digital Disk Recorder

dpsVelocity – 2D
The combination of hardware and software is what makes this real-time NLE system the ultimate choice. The **dpsVelocity – 2D** delivers faster, sharper, and easier edits at an affordable price. This real-time editing solution includes the dpsReality card, dpsReality player, audio/video breakout box cable, dpsVelocity NLE software, and so much more **2699.00**

dpsVelocity FASTER, SHARPER, EASIER EDITS

dpsVelocity – 3D **4499.00**
Edit Bay/Video Action 6.3 ..... **429.00**
Spark Plus/Video Action 6.3 ... **459.00**

DPS

Information: 262-548-8120
Fax 262-548-8130
Tech Support / RMAs:
262-548-8159, 12-4 pm Mon.-Fri.
Safe Harbor Computers
W226 N900 Eastmound Dr.,
Waukesha, WI 53186

arted with ve Modeler

by Sam J Bowling

When I begin working on a new model, some of the most important steps start before I ever open LightWave's Modeler. In fact, a good deal of the planning and designing of many of my models is done away from the computer, using only pencil and paper.

If you're working on a realistic human character, keeping things in proportion is absolutely essential. Most people see hundreds of people a day through traveling or television, and they become extremely sensitive to any deviations from the standard form. While we may not always know exactly what is wrong, we almost always notice when something isn't right. A little thing like putting the eyes too close together can ruin an otherwise perfect model. Mistakes like this are made all too often, and not always by a novice modeler. When you are trying to build a model that mimics life, it is important to have a good understanding of the subject you are trying to create.

In creating a creature of your own design, drawing skills come in extremely handy. While you don't have to be Leonardo da Vinci, a little drawing skill can save you hours of trial and error in Modeler and drawing becomes even more important when you are working with other artists. As the saying goes, a picture is worth a thousand words. Being able to sketch out ideas for 3D models is an invaluable talent and the better you can draw, the better you can organize your ideas. In an ideal situation, taking a few years of art classes would be the best thing to do. While I highly recommend this, it isn't always possible. The other possibility would be to buy some books on the subject; they are always there when you need them and can be much cheaper than a college course. Some of my favorite drawing books are those by Walter Foster. They deal with drawing many subject from the human figure to horses to cartoons and are really well written. Another excellent book is *The Figure* by Walt Reed. It covers just about everything you could want to know about the human figure for artistic purposes. You should be able to find these and other good books on drawing at your local bookstore or arts and crafts stores.

When I'm thinking up a new scene or model I'll grab a pencil and paper and go sit in a comfortable chair for a while and sketch out some different views and make some notes of interesting details. I'll sometimes even create a history for the subject to help work out the details. Where did it come from? What did it eat? Was it hunter or hunted? Remember, the better you know your subject, the better you will be able to recreate it.

Some knowledge of drafting (engineering drawing) and anatomy can really enhance the look or your modeling work. For example, you should know that movement in a robot often uses hydraulics that move the limbs by pushing, as contrasted with movement in a human by muscles which can only pull. Getting the engineering or biomechanical design right can enhance not only the look, but the believability of your work. Things like this are overlooked in many science fiction characters. I'm constantly seeing creatures that have no muscle mass, long skinny limbs and a waist the size of my wrist portrayed doing things that require enormous amounts of strength. In all likelihood a creature like this would live on a low-gravity planet and be much weaker than your average human.

When you're ready to work on the design of your model you may want a source for some reference material or just some ideas. Where would you find this? If you have a computer there's good chance you have an Internet connection. If you don't have an Internet connection you're missing out on an incredible amount of information and the possibility of getting advice from some of the leading experts in their fields. The first place to look for information on the Internet should be a search engine. Actually, it should be on many search engines. Each search engine works differently and may retrieve different information. Just because you can't find what you want on one doesn't mean the information is not out there. Another important thing to remember is that anyone can have a web site, and the information is not always accurate. It's important to check the integrity and accuracy of the site before relying on information you find there. I recently downloaded a dinosaur image from a web site to use as a reference image, and it turned out that the image was mislabeled. While this cost me some time it could have been a lot worse. If you were to download and use a copyrighted image in your own work, you could find yourself in court.

So you've gone through all the search engines and couldn't find what you need. The next step would be to search through news groups. Deja.com keeps a record of most news groups and enables you to search them by keywords, much like an Internet search engine. While this still may not get you the exact information you are looking for, it can help you find a news group that deals with the subject you are interested in. Go there and ask if anyone knows of any web sites dealing with the subject you're researching. You'll usually get an answer within a day or two if you ask nicely and anyone knows the answer, although I've sometimes received additional replies weeks and even months later. Many times the replies I've gotten will reference a book or magazine. If you have the money, I would suggest buying the magazine or book if you will be doing a lot of work related to this subject and need a constant reference. Sometimes the books mentioned will be out of print, difficult to find or you just can't afford to buy it. In a situation like this the library is your best bet. If they don't have the book, ask them if they can get a copy through Inter Library Loan, ILL. More often than not they will be able to locate a copy, although you may have to wait a few weeks. Finally, for additional reference don't forget the movies. Movies have been made covering nearly every subject possible. This is great reference for studying how your subject moves or getting inspiration for an original creation. Be sure not to limit your reference movies to big Hollywood films; documentaries cover nearly everything that lives or has lived on the earth and are full of important information on many subjects.

You need as much information as you can get, starting with measurements. If you're building an animal, say a dinosaur, you'll not only need to know how tall it stands, but how long it is and how wide it is. You'll need to break this down further. How long or wide is its head? What about its legs? How long is its tail? How big are its teeth? Measurements are only the start. You'll need to know how many toes and fingers it has, how many teeth? What color is it? What is its skin texture like? Does the skin have large folds like an elephant, or is it smooth and tight fitting like on a young human? How are the muscles attached and how is their appearance affected by movement? There is much more information that can be useful, but I don't have room to list everything here and it all depends on what your subject is.

One thing that can really cripple your model is not paying attention to how your subject moves. Movement adds a whole new twist on how things should be modeled. If you're modeling a house that will be dancing and twisting across the screen you'll definitely need to model it differently from a house that will not be moving. When working on human characters, you should be aware of the range of motion of the limbs and joints. This is a place where some knowledge of anatomy would come in extremely handy. A good example of this would be twisting a hand. I've seen a few animations where the hand would rotate at the wrist. This would be impossible for a real person to do. Twisting your hand involves movement down your entire forearm. Not knowing this or assuming your audience wouldn't notice something like this could be a big mistake. I can't stress strongly enough how important a thorough knowledge of your subject is.

How a creature moves and where its skin bends and stretches have a big impact on how it should be modeled. A human has soft and flexible skin that can easily bend and flex, while an elephant has extremely thick skin that tends to have wrinkles where it flexes a lot. Once again information on the muscle and bone structure of your subject come in extremely handy. Everything here is important in creating a believable character that will look good when animated. You don't want to bump map a large fold of skin onto your dinosaur model any more than you would want to model every single scale (or feather) on it.

So you want to create your own unique creature. Is all this information useless? Of course not. Just about every science fiction creature ever created was based on one or more existing creatures. It could have a snout like a wolf, eyes like a snail, skin like an crocodile, arms and legs like a human and claws like a dinosaur. You can even use the resources mentioned earlier to find inspiration when you have no idea at all what you want to make. Try a search on Science Fiction or Fantasy or Dinosaur. You could also go through these sections in your book store, library or video rental store. If none of this gives you any real inspiration, try your back yard. I've found some of the most fascinating and alien looking creatures under old rocks and stumps in my back yard. Be careful though, some of them bite.

Once you have your reference books, movies, images from the Internet and sketches, the next step would be to make some background images. These are another important part of getting good results. They can be anything from line drawings to photos. I prefer line drawings because they use less memory and redraw quicker in LightWave. It's also a lot easier to see the edges of things in LightWave's four color dithered view ports.

One of the most important factors in creating backdrop images is that everything needs to line up from one image to the next. If the eye in the image for the side view doesn't line up with the eye in the image for the front or top views, it's going to cause problems when you're modeling. When I make my backdrop images I'll start by just sketching out one view, usually the one with the most detail since this will be the most difficult to draw. Try to keep the details to the minimum necessary. If I was creating a fish, for example, I would draw the outline of the body, eyes, mouth, gills and fins, but I wouldn't draw in the scales since I won't be modeling them and they would just add confusion to the background image. These little details are usually best left for the surfacing stages. Once I have the first image done I'll draw thin construction lines extending horizontally and vertically from all the major points to help me draw my other views, Figure 1.

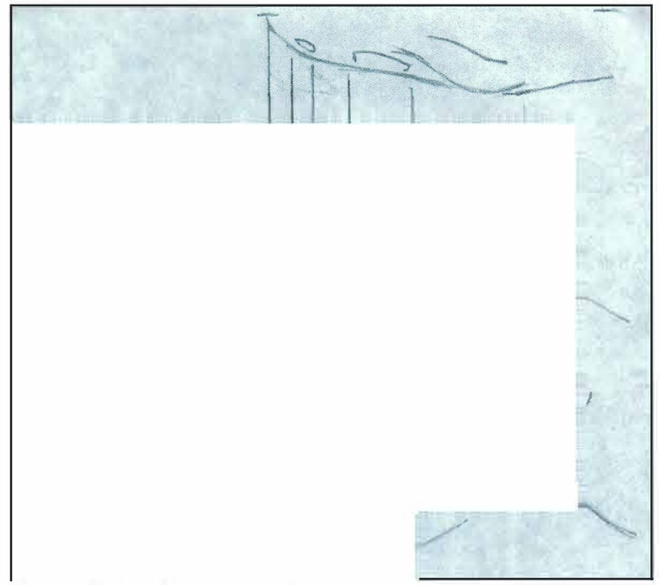

Figure 1: Sketch with construction lines.

Once this is done I'll bring the sketch or reference photos into Photoshop and trace it into a new layer to clean up the drawings and darken the lines, Figure 2, for better visibility in Modeler.

You'll notice that the top and front views are only half of the image. This is because the two halves will be identical so I only really need to build one half in Modeler, then mirror it and merge points to complete the object. Before the background image is ready we have a few more things to do to it. First, copy each of the drawings to a separate image file, making sure that we get right up to the edges of the drawing when we copy, Figure 3.

Double the size of the half images by enlarging the canvas; keep the original image oriented to the correct side of the canvas when doing this, Figure 4.

You can save a lot of time in setting up the backdrop images by setting them all to the exact same size. Determine which image has the largest pixel size, in either width or height. I found that my side image was about 498 pixels wide, so I rounded up and resized all the images' canvas sizes to 500 in both X and Y axes. This produced perfectly square background images with the center of the half-images lined up perfectly along the center of the canvas. One last thing before saving these images is to rotate the top view image -90 degrees so it has the proper orientation.

With the background images finished it's time to set up Modeler. One thing that's

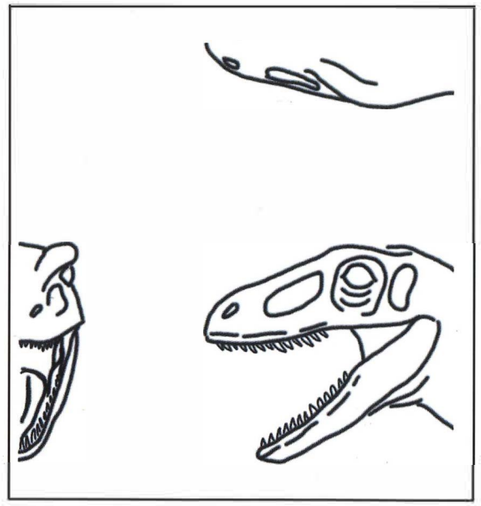

Figure 2: Scanned sketch with lines darkened for better contrast in Modeler.

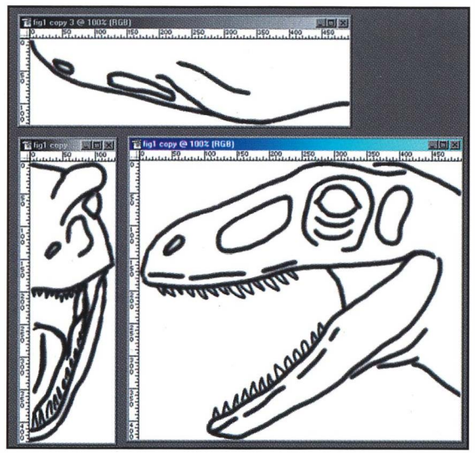

Figure 3: Sketch split up into separate image file for each view. [fig3.tif]

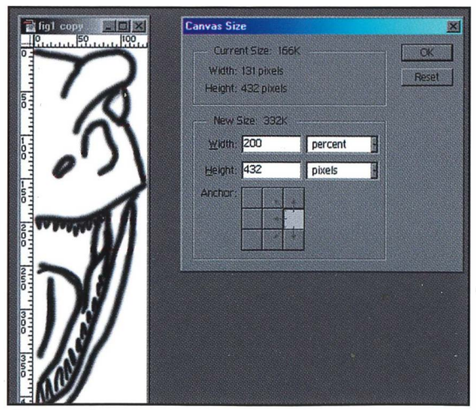

Figure 4: Doubling the width of a half-image.

important when modeling is to have easy access to different viewports and display options. LightWave enables you to configure the numeric key pad to control these settings. I set the 1, 3, 7, 9 and 5 key to quickly switch to the full screen face view, left view, top view, perspective view and quad view. Note how these key shortcuts mimic the position of the viewport they are assigned to. The 2, 4, 6 and 8 keys switch between Wire, Frontface, OpenGL sketch and OpenGL smooth shaded perspective display modes.

Setting these up couldn't get much easier. For the first viewport shortcut key setup, select the center of the quad viewport and drag it all the way up to the top right of the screen. Press the CTRL key and the 1 key on the numeric pad at the same time. The Save View Parameters window will pop up. Deselect everything except quad shape and click OK. Repeat this process for the other four views described above. It's important to make sure that you have the Full View windows opened as far as they go when creating these shortcut keys or some functions like Fit Selected (a key) will not work as expected.

Setting up the shortcut keys for display modes is similar to setting up the viewport keys. Press the d key or select Options from the display menu, then change to the mode you want the first display key to be set to. In this case I selected Wire, pressed CTRL + numeric pad 2 then deselected everything except Preview Type from the Save View Parameters. Go ahead and set up the rest of the numeric keypad keys as I have above, or if you feel you have a better layout go with that. Just make sure you have a setup that is easy to remember. The less you have to think about these things the faster and easier you will be able to work on your models.

While I'm talking about shortcut keys, I feel it's a good idea to stay pretty much with the defaults for Modeler, just to prevent confusion when working on a different machine. I have, however, added a few extra shortcut keys, including F2 for Center, F3 for Center 1D and F5 for Add Points. To setup these shortcuts, go to Edit Keyboard Mapping (F1 on Windows machines) and select the shortcut key you want to use on the left and the plug-in you want to use on the right.

While I think LightWave is one of the best 3D modelers available, there are some free plug-ins that make modeling even easier. The most important one for me is Plane Meta Splitter, which has been around since at least LightWave 5.0. This plug-in splits a row of quad polygons along the 'even' or 'odd' edges. This allows a row of polygons that bend and twist to be split in half easily, and ensures that the polygons will remain quads. It's great in places where the Knife tool won't work. The Intel version of this plug-in can be found at: ftp://ftp.newtek.com/pub/NewTek/Lightwave/Plug-ins/Intel/PlaneMS.zip

Another excellent tool is the L-script selzero.ls. This can save a lot of time when point modeling. For example, sometimes I want to copy a row of points and move them, then use the new points to build some new polygons. The way I used to do this was to select the points, copy them to a new layer, move them so they weren't over the old points, then cut and paste them back into the original layer. With selzero, all I need to do is select the points and run the L-script. To make things easier you can add it to your user menu through Objects|Prefs|Edit commands. Once you have it in your user menu you can create a shortcut key to it using the same method described earlier. The selzero L-script is available at http://www.best.com/~sfnaqvi/lscript.html.

I hope you find these tips useful in your own modeling. If you have any questions or suggestions, I'd like to hear from you.

Sam J. Bowling runs Third Dimension Illustration. You can email him at sbowling@home.com, or visit his website at http://www.members.home.net/sbowling.

We're STILL The Best!

TutorVision 3-D CD

TutorVision3-D CD Volume 4
Beginners Guide to LightWave & Inspire3D!
The ultimate set of introductory tutorials for new LightWave and Inspire users! This CD will demonstrate many basic concepts and techniques that new users will need to know before they can get the most out of these incredibly powerful animation tools! All of the most Frequently Asked Questions are answered in easy to understand video and text tutorials that can easily be viewed as you work on the computer!
Available NOW for ONLY $49.95!

TutorVision3-D CD Volume 3
MORE THAN 6 and a half HOURS of video!! Cool NEW techniques for LightWave & Aura! Topics covered include advanced motion control through the use of null objects, beginner and intermediate compositing techniques, tips for using SkyTracer and Steamer effectively, tapping into the power of HyperVoxels & ParticleStorm, and more! Also contains a number of product reviews in video format and a whole new batch of images from the LW mailing list contest!
Available NOW for ONLY $49.95!

TutorVision3-D CD Volume 2
The most popular and most comprehensive tutorial CD ever created to teach Character Animation with LightWave! This 2-CD set has more than 6 HOURS of video tutorials and will teach you everything you need to know to start doing Character Animation with LightWave 3D version 5.5 or greater! Includes an interview with Stephen Bailey, dozens of images from other LightWave users, a dozen ready-to-render objects, a dozen texture maps, and MUCH more!
Available NOW for ONLY $49.95!

TutorVision3-D CD Volume 1
The FIRST video tutorial CD ever released for LightWave! This disc contains nearly 4 HOURS of video tutorials that focus on doing many science fiction effects as well as more practical every-day effects. The CD content was created with v5.0, but all techniques work fine with newer versions! Includes a video interview with Alan Chan, dozens of images from other LightWave users, a dozen ready-to-render objects, a dozen texture maps, and MUCH more!

LightWave Mailing List Archive Search Engine
An easy-to-use browser and search engine allows you to quickly tap into the information found in more than 50,000 messages from the LightWave Mailing List!
Specially licensed search-engine technology gives you the ability to easily find answers to any number of questions discussed on the LightWave mailing list in 1997 and 1998!
(PLEASE NOTE: This CD is only available from Event Horizon Studios!)
Available NOW for ONLY $19.95!

Any 2 Volumes for only $90.00!
Any 3 Volumes for only $130.00!
All 4 Volumes for only $165.00!
All 5 of our CDs for only $175.00!!
Order before January 31st, 2000 and get Priority Mail shipping for FREE!!

Here's what some of our customers have to say:

I loved this CD and it has been very useful.
- Russ Baker <rbaker@worthington.com>

These CDs are a "must have" for any LW user...beginners especially! They don't replace the manuals but they fill many of the gaps left by the manuals. I wish I bought these CDs sooner!
- Ken Kepf <kmk14@earthlink.net>

Your IK video really answered a lot of questions about Inverse Kinematics!
- Dave Hardman <scbvideos@yahoo.com>

TutorVision3-D is awesome! Mind you I haven't tried anything else to teach me bones...Probably because TV3D Volume 2 taught me all I needed to know. I Highly recommend it!
- Jonathan Cardno <cardno@clear.net.nz>

I have a copy of TutorVision CD Volume 2 and I must say it's a very good introduction in allot of tricks and ways to do character animation. The basics are very well explained and after you understand those you can go as far as your imagination reaches...
- Frank Aalbers (Station X Studios)

Try our FREE Sampler CD!

TutorVision3-D - Sampler Disc
This FREE SAMPLER CD introduces you to the TutorVision3-D series of CD-based video tutorials. This CD contains 1 complete video and text tutorial from each of the four currently available volumes of TutorVision! It also contains 2 brand new video tutorials not found on any other TV3D volume! That's 10 video and text tutorials for FREE!

So how do you get your copy of this FREE Sampler CD? Just get your mailing address to us via e-mail, phone, or use the special form on our web-site, and we'll mail the Sampler CD to you right away, with no obligation to buy anything else, no hidden fees, and no strings attached!
Offer expires January 31st, 2000

Visit us on the World Wide Web for FREE Demos, Program Updates, and LOTS of other Cool Stuff!

Recommended System Configurations
PC: Pentium, 16mb RAM, audio card, 4x CD-Rom, Win95/WinNT
DEC Alpha: FX!32, 16mb RAM, audio card, 4x CD-Rom, WinNT
PowerMac: 16mb RAM, 4x CD-Rom, System 7.5 or greater

Event Horizon Studios
ORDER NOW!!
Phone: 518-756-6850
FAX: 518-756-6836
E-Mail: info@tv3d.com
Web-Site: http://www.tv3d.com
We Accept:

Most of my plate matching tips are not actually lighting tips. They are a means to an end. A good portion of this article covers LightWave math and how different lighting variables work together. I hope that this guide will help you choose the quickest and most calculative means for lighting to plate photography. Most of my techniques are designed to eliminate guesswork and make it easier for you to debug your lighting scenarios. Matching BG plates requires you to simultaneously tweak your surfaces, lighting, and atmospheric effects. The techniques described in this article will give you better control and easier access to getting what you want out of LightWave's renderer.

When you start applying textures and setting up your lights and fog, you are making an equation or flow chart for LightWave to compute. This result of this computation is your rendered image. It's very important to remember this fact: 3D programs are still just programs. They take (this number) and multiply it by (this), then run (this) using (this) operation, and so on. In LightWave we have a lot of variables to work with. Some are useful, some are not. The fewer variables we have, the easier and more manageable our tasks will be to get the results we want. We still need many variables to achieve photorealism, but we should eliminate any redundant or unnecessary variables.

Here are the main points of my lighting approach:

· Set up your camera lens, position, and rotation using on-set measurements; if you don't have that, use your best judgement.

· Model a portion of your plate, using proper scale, to be used for lighting tests and shadows.

· Make sure the areas that you choose to model are ones that will be near your CG elements.

· Extract simple textures from the plate photography, making sure to rid them of any lighting effects.

· Position your model to the appropriate distance from the camera, making sure to check the top view.

· Set up the position and angle of your key light using a 255-255-255 light with 100% intensity.

· If you are going to use colored lighting, color correct your textures to their original or native color.

· Adjust your light colors by sampling colors from, and analyzing, your plate photography.

· Keep all surfaces' diffuse parameter set to 100% to eliminate the diffuse variable from your lighting equations.

You should keep Photoshop or your favorite compositing program running in the background with a frame of your plate photography in it. In Photoshop, make a grayscale version of your BG. This will help you to more easily judge shadows and diffuse shading. Experiment with auto-level and levels to crank up the contrast and eliminate fog effects. Now you have a clear picture of how light is hitting different areas of your plate. You can see how this has been used with the "Fire pit" plate in Figure 1. For color sampling, it's a good idea to blur your original plate to get rid of grain, and use this image for a more average color sample.

I highly recommend using the QV 3.0 image viewing plug-in for your test renders. It lets you keep layers of test renders, add notes, and reads out RGB color info when you have your mouse over an image. You can also use it to save images or copy them directly into Photoshop. If you put your BG plate into LightWave's Effects Panel's Foreground Image, LightWave will skip its render and just output your BG image. You can keep this in a QV layer for reference and to compare with your test renders.

LightWave works from a solid system of mathematics. If you shine a standard 100% light directly at a surface set to 100% diffuse, you will get the exact color values as your surface color or color texture. In other words, what you see in Photoshop is what you get; it's WYSIWYG. A light color of 255-255-255 at 50% intensity is the same as a light color half that, or 127-127-127, at 100% intensity. This is a major key to getting your colors and intensities correct and eliminating lighting variables. There is no difference between making a light 50% or making your surface 50% diffuse. There's also no difference between halving your surface's color versus using a diffuse setting of 50%. Take a look at Figures 2A and 2B; the rendered colors are identical. Figure 2A has a surface color of 190-128-64 with a diffuse setting of 70%. Taking 70% of the surface color values 190-128-64 leaves us with 133-90-45, the color values for Figure 2B. When rendered with a diffuse setting of 100%, we get the exact same results as in Figure 2A.

Diffuse properties in LightWave are usually unnecessary, adding a redundant variable to your efforts. As noted before, you can get the same exact same effect by multiplying a diffuse value by your color values and leaving your diffuse at 100%. For texture maps, you can do this yourself in Photoshop by layering your diffuse maps on top of your color map and setting your diffuse layer to multiply, as shown in Figure 3. A unified color and diffuse map multiplied in Photoshop will look the same as separate color and diffuse maps multiplied in LightWave. The idea of using a diffuse map for dark and light areas is fine; it will be a lot easier for you to tweak if you multiply it yourself in Photoshop. This way you will know from looking at your image map exactly what colors and values you are going to get from your render. In LightWave 5.6 you can't change the mathematical operation (add, multiply, etc.) used to blend a texture layer, which means that adding a fractal noise of black to the color channel is not the same as multiplying fractal noise as you would by using the diffuse channel.

Always use a key light of 100%, and try not to use colored lights. This is especially important when your scene has a lot of color variation or daytime lighting. Obviously, use colored lights if your shot requires it, such as matching a bright orange fire pit that your character needs to walk by. If you need multiple colored specularity, consider using a lighting plug-in or use the clone light/no diffuse/no specular technique. If you haven't purchased any lighting plug-ins, I highly recommend them. They give you much more control for achieving photorealism. You can edit the diffuse termination, diffuse transmission, specularity, and include and exclude lights. Without lighting plug-ins, you'll have to jump through more lighting hoops to get the same effects.

Why do I recommend using 255-255-255 colored lights? This goes back to LightWave math and value relations. Colored lights are another variable to tweak and add to your equation. In most cases, you don't want to use colored lights in your initial lighting setup. Slight tweaking of light color values in later stages can really help you out. If you use colored lighting, you'll want to adjust it simultaneously with your surfaces. If you make a light slightly orange, every surface is going to be affected with slight orange. Take a look at Figure 4. I have a simple model with a simple texture for the face; I chose this area of the background because it's close to my CG elements. Notice that the light shining on the actor's left side is colored a slight green. I did this to match the background plate, since the light to that side really was a slight green. Changing the light color was one of my last steps in lighting this scene. Get the intensity of the light right first, and only then adjust its color.

A lot of lighting artists use special "kick" lights with intensities of 200%+ for getting really bright, specular kicks or highlights. This same effect can be achieved by cranking your surface specularity or diffuse to over 200 percent. Yes, in the relative nature of LightWave it will render exactly the same, as shown in Figures 5a and 5b.

For color correcting, always try to get your colors right through your color maps and surface color properties. When lit with white light, these maps and surfaces will give you the colors you are looking for. Your render will look just like your texture maps do in Photoshop. This technique has been used to great extent in digital matte paintings and set extensions. Remember, if you turn down the texture opacity in a color texture panel, some of your surface color will show through. Although I recommend getting your colors right within the color map, this is a great place to tweak colors. Since your texture is dissolved with your surface color, this can sort of be used as a fog per surface (except that LightWave's fog is luminous). Using LightWave's colored fog will add yet another variable to your pursuits. If you need fog, you'll want to adjust it simultaneously with your lighting and surfacing.

You must be able to adjust all of your surfaces for every shot. This is most important when finalizing surfacing. Since LightWave doesn't have the ability to save separate surface scenes, I do this by loading all my geometry into Modeler, polygon reducing all of it, and saving it out as a surface object. I usually save one for every scene, naming the object by the matching scene name. You'll probably want to do this in the final stages, after all texture maps are applied, and all of your surface selections have been assigned. Otherwise, you will have to painstakingly update all your surface objects. When making your surface object, remember that all you need is one polygon per surface. I'll often just hit w, select some of the more polygon-populated surfaces, and cut all but one polygon away. However, this can mess up automatic sizing, which is another reason hold off on this process until after you have your textures applied. Finally, always make sure your surface object is loaded into the scene last to insure that it overwrites all the other surfaces. This enables you to make surface adjustments to a single object in Modeler and have them automatically applied to all the objects in the scene.

The ambient light setting in LightWave needs to be treated just like any other light. Often, it's better to use more lights for your ambient needs, as this can help bring out contrast. Use this with caution, as CG images often have too much contrast. However, I once used an ambient intensity of 35% to achieve photorealism, so this can certainly vary. It really depends on what you are color sampling from your background plate. Pick the darkest area or shadow and sample the color values. If it's above 0-0-0, then there is ambient light in the scene. Is it ambient like LightWave's bland ambient, or is there more contrast in these dark areas? For more contrast, you will probably want to add light, aimed in a different direction of your key, instead of using LightWave's ambient.

Figure 1: Grayscale version of background plate.

Figure 2: A, above, with brighter surface color and lower diffuse value. B, to right, with darker surface color and higher diffuse value.

Figure 3: Photoshop used to multiply a color image by a diffuse image.

Figure 4: LightWave lighting matched to a background plate.

This is certainly true in the "fire pit" example, because even the really dark areas still have detail, as shown in Figure 1. You could clone your key light, lower its intensity, and flip it 180 degrees, which will bring in light from the opposite direction. This can be a good place to start for your "ambient" lights.

When light and surface tweaking, always refer to your color sampling. I usually keep a pad and pen handy or keep WordPad open to jot down color values. I'll sample a high color, a mid color, and a low color from the background and write all of these down. Then I'll sample the same areas of my current test render. This makes it easier to see what variables need to be cranked up or down. One nice thing about using white lights is that all you have to do is tweak the surface's color to match. If you are using colored lights and/or colored fog, you'll have a few more variables to tweak and put into your equation. Remember that all LightWave is doing when it renders is computing your equation. There are a lot of options, so remember to always choose the most exacting one. I hope that this guide will give you a good clue as to when to use what technique, and why. In Part II, we will go through a real-world example of lighting matching using these techniques.

Matt Morgan is a senior animator at Threshold Digital Research Labs. You can email him at violator@thethreshold.net.

150% Light,

Figure 5: A, above, with brighter light and ambient values and lower diffuse and specular values. B, below, with lower light and ambient and higher diffuse and specular values.

sample and the compressed DigiSuite LE sample. DigiSuite LE was set to its highest quality setting. In these samples, the slight compression is not visible at all, but the compression makes it much easier to store and manage footage. Additionally, if you ever need to capture an unusually large amount of material, you can increase the compression ratio and double or triple your system's storage capacity without even turning a screwdriver.

Despite the problems I ran into, the Video Toaster does have a lot to offer. There were many things about it that I liked. At only $2995, NewTek has significantly undercut the uncompressed competition. Any other system that offers uncompressed capture will cost you a minimum of three times as much. The image quality the Video Toaster NT delivers is stunning. Another nice advantage to this pristine signal is that the Toaster can double as a great replacement for your digital disk recorder. The inclusion of a built-in software waveform monitor and vectorscope is a wonderful touch, and I really missed not having this in the DigiSuite LE. Aura, LightWave VT and Speed Razor alone without any hardware would be a fair deal at this price. Toaster NT does give the user a lot of bang for their buck. And of course, in the NewTek tradition, if you own a previous Toaster system or other NewTek products you may qualify for a sizeable discount. Contact NewTek for details.

If you are considering a Toaster, you should definitely put it in a system that has been qualified by NewTek. I have seen Toasters work better than mine as long as you adhere to system recommendations and stay within the suite of tools NewTek provides. For many users these tools will be more than adequate for production. While the original Toaster was truly revolutionary, the new Toaster seems much more evolutionary. As Sade would say, "It's never as good as the first time."

After using both of these products have to say that I prefer the DigiSuite LE. Unless your applications demand totally uncompressed video, I cannot see a reason not to use the DigiSuite LE. In addition to its built-in features, users can opt to purchase a breakout box for audio and video as well as the Pinnacle Genie real-time 3D DVE board. Matrox has also announced that in the first quarter of 2000 it will introduce its own (likely less expensive) 3D DVE processor as well as a codec board that will add real time DV and MPEG-2 for DVD encoding and transcoding. DigiSuite LE is a far more mature, more complete package than the current Toaster. It is NewTek that should be running for the R&D lab. It doesn't take long to have your lunch eaten in this industry.

As a Mac user I would not hold my breath for a Mac Toaster, or a Mac DigiSuite LE for that matter. Apple's near death two years ago derailed a lot of manufacturers' development strategies. With key players like Avid and Media 100 making the switch to NT, it seems less likely that someone new would enter the Mac market. Mac users looking for quality and performance in the range of these products need to look to boards like the Targa 2000 DTX or RTX. If NewTek can squash the big ugly bugs crawling all over this latest Toaster, they will have a strong product on their hands. Until then I will be looking to other systems.

Jason Goodman is the CEO of NYC based Multimedia Software Design. You can email him at jason@msd nyc.com or visit their website, www.msd nyc.com.

Figure 3: While both cards are single slot devices, the double-thick DigiSuite LE (bottom) contains a lot of on board hardware to handle realtime 1.3:1 compression. The uncompressed Toaster card relies on software and fast CPU speeds.

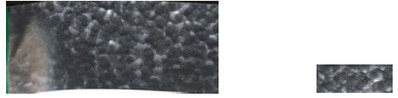

By its very name, LightWave is all about light. The word is even in the software's name, which should give you a clue. Mastery of lighting setups will give you mastery over LightWave more than any other single technique, which is why we are talking about it in this tutorial.

Light in LightWave is handled in a ray-tracing method; that is, light travels from the light source, hits the surface of the object and stops dead. It's like a bug or a BB pellet hitting a soft peach at speed. Ker-thwap. Stopped dead. Unless the surface is a mirror, the light is not reflected in any way. If you want to send light anywhere other than straight onto the object, you have to direct it yourself, and this means you have to cheat the effect of lighting in the real world. Of course there's a good chance that in L[6] all this fake Radiosity talk may become redundant. However, my opinion is that knowing how to do something manually is never wasted space in your brain. You never know when that knowledge will come in useful.

Lights, Camera, more Lights…

The first thing you have to know about is Ambient light. In Layout, load a simple object. In this example I'm using a shape I've created using MetaNURBS and textured in Layout. With the object centered in the screen, render it to a default lighting setup. You can see the model because there is a main light (called the Key light) in the default lighting setup, ready and waiting every time you start LightWave, and the Ambient light setting (in the Lights panel) is set to 25% by default. You can test this. Go into the Lights panel and set the Light intensity to 0% and ramp up the ambient light to around 75%. Close the Lights panel and render the shot again. Your results should look something like Figure 1. This time you can see plainly that the object is lit in a ghostly way by the ambient lighting algorithm. This doesn't respect any modeling that might be on the object. It just makes it visible rather lighting it. Ambient light is set purposely high so that anyone can load an object and see it in the frame without having to light it properly.

Figure 1: Simple object rendered with 75% ambient lighting.

For professional purposes, we need to turn off ambient lighting completely. This gives us better control and makes any lighting setups we may design behave correctly. Go back to the Lights panel and turn Ambient light down to 0%. Turn the main light back up to 100%. Now when we render the shot we can see that all the shadow detail has gone (Figure 2), and that's exactly what we want. The only light comes from the light we have set, the default light pointing at the default angle, which is to the left pointing downwards and inwards. What we have now is the kind of lighting you can expect on a sound stage in a film studio, perfect blackness all around and the light fully controllable.

Nobody (except snapshot photographers) takes pictures with a single light source. Professional photographers fill in the shadow areas of an object, person or scene with other lights or with reflective surfaces that bounce the light softly into shadow areas. The principal fill lights are a back or top light and a bounce or fill light. To add these lights, go to the Lights panel and create two new lights. Rename the first light as Key Light, and rename the second two lights as Back Light and Fill Light. Turn off Ray Traced shadows for both the fill lights. This is to make the fill lights cast no shadows and give them the feeling of diffuse bounce lights. We could also click the No Specular button on each of them, but we'll try that later on.

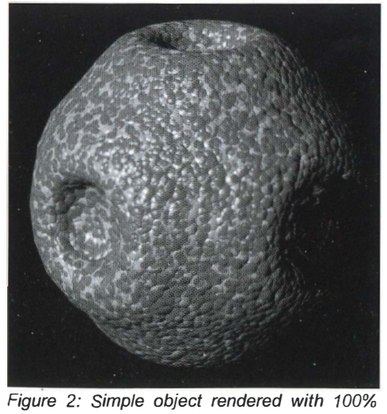

Figure 2: Simple object rendered with 100% default light and 0% ambient lighting.

to Layout and
click the Lights button
in the Edit menu. You are now
ready to set the position of the lights. The lights we have set up are distant lights. This means that their position in the scene is irrelevant. The light from these sources runs parallel through the scene in the direction the light is pointing, regardless of the light's location. It's like a big white laser.

Choose the Back Light, and rotate it 90 degrees so it is pointing down. You may prefer to move the light up along the Y-axis so it's not inside the object; this makes it easier for you to see its orientation. You can either rotate it by hand or you can hit the Numeric key and type in 90 degrees Pitch. Do the same to the Fill Light, but rotate this so it is pointing upward, or -90 degrees Pitch. Now when we render the scene, the object is lit from three different directions. I like to make the fill light a little blue and the back light a little yellow, but that's my perversion. I do this for space shots and interior shots, and although it would seem to be a little bit fake, it's the kind of fake that looks more realistic than reality. (I know that sounds confusing, but can you say it 10 times really fast?) For outdoors, it's better to make the Top or Back light a little blue to match the sky, and the Fill light a little more yellow or green to imitate reflected grass or ground. If you have any artistic training, you know that when you paint cold scenes you add warm shadows and vice versa. I don't know why this works, and I don't want to, but work it does and it's a good rule of thumb.

Although the object is looking a lot more interesting now, it would be better to refine the lighting a little. The Back light needs to point forward and to the left a little. Rotate it by hand or type in a Heading of 35 to 40 degrees and a Pitch of about 130 to 140. The Fill Light needs rotating into the scene a little, and back a bit too. Rotate it by eye or set a heading of about -40 and a Pitch of about -45. Now we have something that looks nicely lit. The fill lights and the Key light attack the object from three angles, and so the object will look pretty from almost any direction.

Imitating Daylight

Obviously, lighting isn't something you can provide too many hard and fast rules for. Every shot needs to be treated as a new instance and experience is the only thing that can guide you. For example, if our object were supposed to be lit by daylight we'd have to provide different angles of fill lighting and the overall lighting values would have to be higher. Interiors can be gloomier. A good rule of thumb is to use your eyes and see where the light should be coming from in the scene and to provide it with an extra light as accurately as you can.

A common daylight scenario might go like this: The sun is very bright so the Key light needs to be very bright, well over the basic 100% light that is provided as a default setting. If you go into the Lights panel and ramp it up to about 150 to 200% this better replicates sunlight and starts to put a bloom onto your highlights. Put the fill lights back to 90 degrees for the Back light and -90 degrees for the Fill Light. Make the Fill Light slightly yellow and the Back light slightly blue. Make changes to the intensity values for the fill lights, making the Back light 150% and the Fill Light 100%. Click the No Specular button for both fill lights.

What we have done is to make the fill lights better mimic the kind of light you get outside. A lot of light comes from the sky that is not directly from the sun. A little reflected light coming from the ground, which you get even if it's a fairly dark surface. But because diffuse light sources like the sky and the ground are not direct lights, you will hardly get any highlights from the light coming off the ground. Therefore, a No Specular setting kills the highlights and makes the fill lights more even and diffuse.

Figure 3: Mimicking daylight.

Try this: Make the Backdrop color a light blue, 78,112,150, and if you're feeling brave you can experiment with a little backdrop fog to imitate the particles in the air diffusing the shadow detail. (We'll be going into this in more detail when we talk about compositing in a future article.) Render the scene and you can see that we are really starting the change the look of the thing, Figure 3. It looks like we are outdoors on a sunny day. The darker areas of the shape are filled with light of the appropriate brightness and color, and the sunlight is sufficiently bright.

Creating a Gray Day

Harder to do is a gray, overcast day. Of course there are no sharp shadows on a cloudy day, but there is a lot of light around still. An object that is shiny still looks shiny even on a gray day, but the main source of light is now the sky, filled with clouds and illuminated by the sun that is above them.

Make the Backdrop color a little more gray, and then adjust the lights as follows: a Key light, which is above the object with Intensity in the Lights panel at 100%. Rotate the light so it is pointing straight down, a Pitch value of 90 degrees. Turn off Shadow Type. (Diffuse light casts no shadows.) Next we adjust our fill light. Drop the intensity down to around 25-30%.

The next step is to add five fill lights, one pointing up, one from the left, one from the right, and another pair pointing forwards and backwards. The idea is that to fill the object with diffuse light we have to light it from all six sides; above, we have the Key light and from the sides, back, front and from beneath we have the fill lights. With this in mind let's clone the fill light and position the clones around object. Select the Fill Light and click

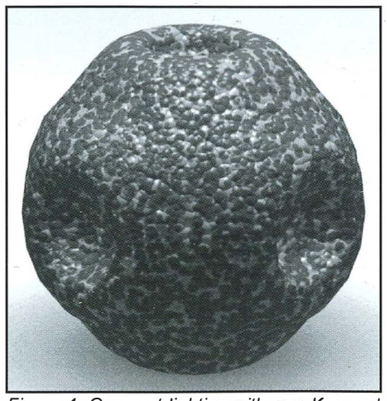

Figure 4: Overcast lighting with one Key and five Fill lights.

Clone Light. Make four clones. Rotate them so that they point at the object from the side, back and front and leave the original underneath. We won't need our Back light, so remove it using Clear Light in the Lights panel.

Now when we render the scene the light is splayed out all over the object from all sides creating the feeling of the kind of diffuse light you get on gray days. Obviously if the object is resting on the ground then you will need to create a mild shadow under it using very blurry shadow mapping, Figure 4. A more costly option is to use Area Lights. However, these need to be antialiased to hell and back to stop looking speckled.

Other Ideas

Other things you might like to experiment with are studio lighting setups. How do you imitate the lighting in a studio with all its reflectors and soft window lights? In a studio we have pitch-blackness and we can do whatever we like with subject. When you are going out of your way to imitate slightly stylized lighting, LightWave gives you all the tool s to make this possible.

Use colored lights to create a mood. In our initial example we had a dark blue fill light, a classic photographic ploy. Try experimenting with different colored lights, how they clash or combine and how they affect the finished result.

Professional photographers go out of their way to reduce lens flare and lens reflections, and yet some LightWave artists seem to be heading in the opposite direction. Lens flares should be used very sparingly, if at all. Lens flares really have no place in a realistic lighting setup unless they are used so subtly that nobody would guess they are lens flares.

Finally, experiment with Shadow mapping. It's a costly exercise in terms of RAM usage, but shadow mapped light sources produce shadows that are lovely and soft. If you are using spotlights with shadow mapping, tweak the edge softness levels and be sure to fill in the shadows with a diffuse fill light of some sort.

So don't just cast a light - fill it!

Phil South is a freelance animator and web designer based in the West Country of England. He can be reached via his web site on: www.ii3d.freeserve.co.uk.

Take One

by Lee Stranahan

Four years ago, on my thirtieth birthday, I got a phone call very early in the morning. I had been expecting a nice quiet day at home with my family. I figured I'd do the only thing I really do to celebrate my birthday every year, which is to have my 'birthday meal,' a homemade version of Steak Milanese with a side of highly seasoned buttered corn. That was my vague little plan, but the phone call suggested that I go up to Portland, Oregon right away. So, within a few minutes of waking up from a very sound sleep, I gathered up a bunch of computer files and I was on my way to the airport.

When I arrived in Portland, I was picked up at the airport and rushed off to my destination. I still hadn't eaten any breakfast, so I asked if we could stop for food. So, I sat there in the drive-through lane of a McDonalds trying to decide between a Big Mac and a Quarter Pounder for my birthday meal and I thought to myself, 'Welcome to the exciting world of onscreens. Happy birthday.'

Onscreens are a production element. They are, literally, any graphic element that is visible on a screen during a shot. Every time the villain tries to hack into the computer system and gets a message that says ACCESS DENIED or every time the Good Guy brings up a web page to find out someone's true identity, you're seeing an Onscreen.

It's important to understand that the Onscreen isn't the computer itself. The computer is considered a prop or a part of the set. The Onscreen is the graphics file itself. You can think of Onscreens as electronic props. Just like real props, they need to be created, taken care of and organized so they are ready to go as soon as the Director calls for them.

Onscreens are also one of the only things that a visual effects artist might create that are needed during production. Usually all of your real work is done in post production. As stressful as Post can be sometimes, it is a much more relaxed process than production. During production, the name of the game is to get the shot done now and I mean RIGHT NOW, because the camera is going to roll in a minute and then it's on to the next shot.

To understand the process, let's start at the beginning, which in Hollywood is always the script. The script is the blueprint for the show, so when the script contains a little fake James Bond sequence like this:

> INT. VILLIAN'S MAIN COMPUTER ROOM - NIGHT
> Our hero BURSTS through the door and FIRES off a round between the eyes of a surprised ENEMY COMPUTER TECH. As the dead man falls to the ground, our hero quickly tries to get his bearings. There are a dozen screens showing the different views of the world's major cities, all being targeted. Then, he sees it - the computer that controls missile itself. The status indicator shows that the missile is armed and ready to fire…in ten seconds! Our hero types in a password, but the system won't let him in - and now there are only five seconds left.

This tells the Director that he's going to need a number of onscreens in this sequence.

Let's take a look at the three major steps needed to take the onscreens from script to screen. First, you need the actual screens themselves: a real computer, TV, monitor, laptop, or whatever. Next, you need the onscreen graphics to play back on the monitor. Then, you need to be able to shoot these elements, often with a special sync box or other device. Finally, any interactive elements need to be controlled on set.

The show I was working on was called *Nowhere Man*, which ran on UPN. At the time, I was doing freelance work for E=MC2, one of the leading companies that provides onscreens and playback graphics for Hollywood. A company like E=MC2 provides a 'one stop shop' by taking care of everything that a production needs to make the onscreen happen. That's important, because even though the three steps we just outlined might sound simple, in reality they can be a lot more complex than you might think.

Using our little script segment as an example, let's do a little breakdown and see what will be required. First, we'll need "a dozen screens." Does that mean twelve actual computers? It might. The decision is mainly a matter of budget and the director's vision for the shot. Many times a script will call for some element, but when it gets down to the real world of production it just isn't practical. Let's assume we do have enough of a budget here, and so we are going to try to have a dozen or so different screens.

So, what kind of screens? Don't just assume that a screen is a screen is a screen. There are black and white monitors, big monitors, small monitors, flat panel displays, displays that pivot and on and on, as seen on the next page. This is where a company like E=MC2 starts to have a big advantage; they have literally a warehouse full of different monitors, TV sets, and other display devices. Whether the script calls for a family in the

1950s watching television in the living room, or the head of the Cybermafia wreaking havoc in the year 2112, E=MC2 probably has the display for the job.

Okay, we've picked a dozen displays, so on to the next step. Each one of those screens is going to need something playing on it, too. In this case the script calls for "different views of the world's major cities, all being targeted." What does that mean, exactly? I wrote it, and I'm not sure myself. This is what I meant when I referred to the script as a 'blueprint'. The script is a plan or a guide, not the entire movie. In many cases, it's much more like an idea scrawled on a cocktail napkin then a blueprint, because blueprints usually are very detailed.

In this case, I can think of a number of ways to show cities being targeted. There could be wireframe globes with blinking dots or scrolling maps with concentric circles or even actual videotape of New York and London with little red target overlays. So, how do we know what graphics to create for this shot? One possibility is to ask the producer and director, but bear in mind that everyone on the production is constantly asking them questions so you'll only have about five minutes to figure out what you are supposed to do. Besides, a lot of the time they will want to hear what you want to do. After all, you're the expert and that's what they are paying you for.

What do you do, then? Often, it's just a balance between art and practical considerations. You want the screens to look good, but you are under a production deadline and there's only so much time you can put into the graphics, especially since none of these dozen screens are the 'hero' screen; they are all essentially set dressing. In this case, it's probably best to have a combination of different types of graphics: some text, some globes, some maps, some video, and so on. That way, each screen has its own 'motivation' and the overall effect will probably be pretty impressive, especially at first glance. With background screens like this, that first glance is usually all that the audience will ever get.

Our script sequence also calls for a 'Hero' screen: the missile countdown screen where the password is entered. This is called a Hero because it's going to be front and center, and the audience is going to get a good look at it. That means that you should probably put as much work into this one screen as you do on the other twelve screens combined. This is going to be our 'money shot' in this sequence, and if the graphic looks cheesy or rushed the whole sequence will come off badly.

This brings up a note about realism in onscreen graphics. You have all seen onscreens that looked so fake that you've burst out laughing and thought 'Who designed this?!' Well, probably the producer or director, but I'll leave that complaint for another day. The point I want to make is that onscreens don't need to look real. It is more important that they fit the overall look of the show, convey the needed story information clearly, and look cool. If you can do all that and make it look like a real computer screen too, more power to you.

How do we create all these graphics? For starters, Photoshop. Then, you'll want some 2D animation capabilities such as Adobe After Effects, eyeon's Digital Fusion, or NewTek's Aura. Of course, you'll probably be using LightWave and it's great to have 3D tools in your onscreen arsenal. Typically, you'll also want some interactive authoring tools like Macromedia's Director software. I've used HTML to great effect as an interactive authoring system for onscreens, as well. Just load all the pages off the local hard drive, not the web.

Speaking of LightWave, onscreens are a good chance to use the Rendering Modes you probably never use, like the Wireframe or Quickshade modes. Also, onscreens are a good time to use the 'solid wireframe' trick as shown in the cow rendering on the next page. This is where you take a LightWave object and change all of its surfaces to one surface name. Then, in the Objects panel under Object Appearance turn on Polygon Edges and set the Edge Color to nice obnoxious green or blue. Finally, turn off all the lights in the scene, including Ambient Intensity, by setting their values to 0 and you have a nice little wireframe view with hidden backfaces. It's an old trick, but a good one. Just get that object rotating around and you get a very popular onscreen look.

With the elements created, the next step is playing them back and recording them in the context of a scene. Whether we're recording on film or video, there are technical issues that make this difficult. If you have ever tried recording your computer screen with your camcorder, you know it's not just a walk in the park to record the signal off of a monitor. Without getting into all the technical issues of sync and interlace and scan rate, let's just say that this is an area that you need to discuss with the director of photography well in advance of the shoot date. If the DP seems confused by the whole issue, let that be a warning to you that you need to expect trouble and it is time to call in an expert.

One reason for not going into too much detail is that the solution varies depending on the equipment you are using. Different video cameras have different ways of synching to different monitors, and film introduces whole new areas of fun to deal with, and when you need to synch up two, three, a dozen, or a hundred monitors you get into all sorts of new areas of difficulty. Again, this is why there are companies who handle this sort of thing.

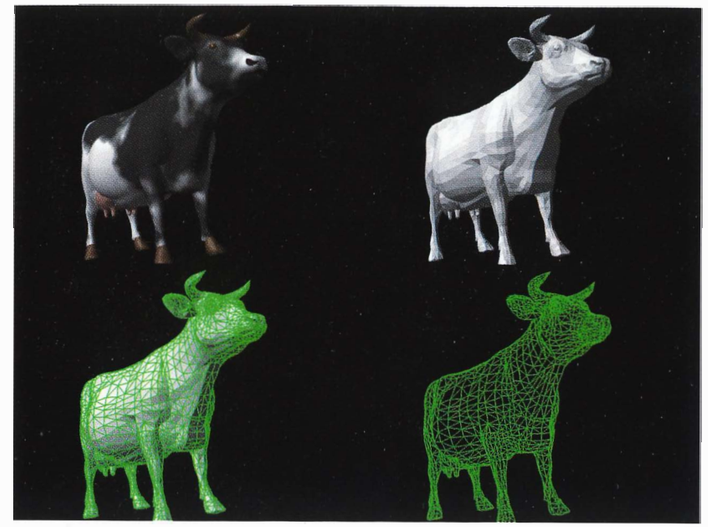

One good general rule, however, is that you avoid many problems by shooting LCD monitors, such as those found on flat panel computer monitors, laptop computers, and those cool little Sony Watchman DV players. LCDs don't 'scan' the way a CRT does, and so they are very easy to shoot.

If an onscreen element doesn't need to be interactive, it can just be run as a QuickTime or AVI movie, usually with a loop of some kind. However, interactivity is often the name of the game: when the good guy types, we want to see letters appear on the onscreen. That doesn't mean the actor is doing the typing, however. Usually, the Director wants the onscreen pre-programmed and foolproof. It's not that actors are all stupid (okay, except for Keanu Reeves). It's that the director wants to do the shot in as few takes as possible, and requiring the actor to type is just one more thing that can go wrong. To avoid those problems, sometimes there is an operator off camera and the actor's keyboard isn't even plugged in. Sometimes, the animation is programmed (usually in Macromedia Director) so that no matter what key the actor hits, the right key appears onscreen. This allows the visible rhythm of the actor's hands to be right, and is a little more realistic.

So, that's a quick overview of the onscreen process. There is a lot of computer graphics and 3D work that doesn't end up as the sexy visual effects shots that we all know and love and want to work on, but end up on monitors in the background. It may not be as sexy, but they do pay you.

Lee Stranahan (www.stranahan.com) is a filmmaker, photographer, visual effects artist and teacher who has taught thousands of people how to get more out of LightWave through seminars, videos and CDs. He lives in Pasadena, CA.

Now shipping!!

Lipservice is a stand-alone facial animation package designed to work with Lightwave 3D and StudioMAX, but also includes support for other packages via Wavefront OBJ.

Features include a powerful built-in modeler, real-time live puppeteering, and one click keyframing. Interpolations include a splined motion path for each vertex (this ensures perfect continuity of motion across keys) and layered delta morphs (for eyeblinks, smiles, frowns, breathing, etc.). Lipservice allows you to scrub your WAV files in real time or step frame-by-frame, marking your keyframes with a single mouse click.

Four years in development and proving, Lipservice is now available for the introductory price of $299 + tax & shipping.

For more information, snazzy animations, or to place an order, check out our website:

http://www.jonalter.com

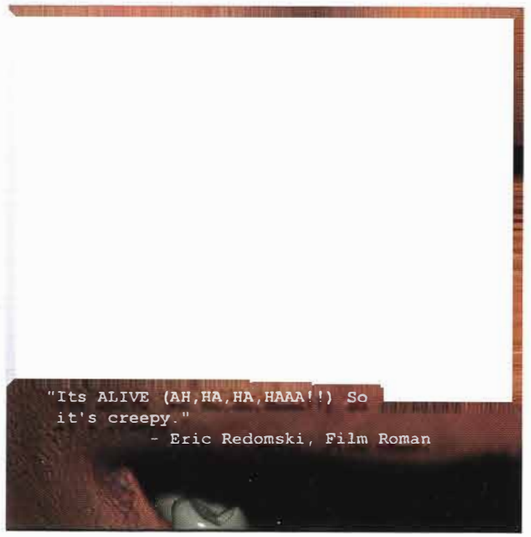

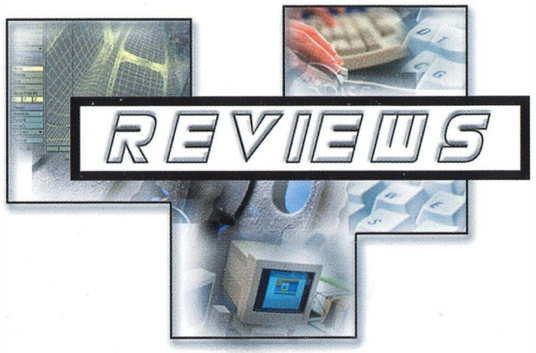

Reviewed this Month:

Software

Surfaces ROM

Dosch Textures CD

Hardware

NVIDIA's GeForce 256

Surfaces ROM

reviewed by Doug Kelly

US $49.95
Distributed by Graphic Detail
4556 South Third St.
Louisville KY 40214
800-265-4041
www.lightrom.com

SurfacesROM is a collection of 300 photographic building materials in categories including marble, stone, pavement, walls, roofs, floors, metals, plants, wood, ground and miscellaneous. The maps are all tileable, and the quality of the tiling is absolutely seamless. If you need to surface a building, tile a roof, add grass to a lawn, pave a walkway with interlocking bricks, fill a driveway with gravel, panel a room in woodgrain or plaster, tile or parquet a floor, or put real bark on your virtual trees, this collection has the maps for you.

If you are doing landscape architecture, interior design, architectural walkthroughs, game environments, or just like to dink around with photographic surfaces, this may be the collection you are looking for. The quality of the high-resolution images is excellent, and the selection is good. The materials photographed are those you would commonly need for household, construction, and landscape surfaces. Most of the color maps also have matching bump, specularity, and reflection maps; those that don't are flat, polished materials that only need general surface settings to simulate properly. The combination of matched maps makes it a no-brainer to set up very nice surfaces.

Each image is available in high resolution (up to 1330x740), low resolution (as low as 160x145), and 300 horizontal pixel resolution (between 131 and 566 vertical resolution). The high resolution set is preferable for the best results, of course, but the RAM conservation of the low resolution set can come in handy. The 300-pixel resolution is especially useful for design work where you or your client wants to see a variety of surfaces for comparison. Because at least one axis of all the maps is the same size, you can swap out surfaces with less worry of a scaling problem.

The HTML thumbnail index makes it easy to peruse the collection, and a nice touch is the inclusion of each image's dimensions. This makes setting Texture Size in LightWave much easier. It would have been even better if the materials were already set up in LightWave surface (SRF) files. As it is, you have to load up to four separate images and tweak their mapping parameters. This is slightly easier with the 300-pixel image set, because both the low- and high-resolution sets have widely varying dimensions. If you purchase this collection, I highly recommend that you spend some time compiling the maps into "standard" surface setting libraries for the scales you usually work in, and save them for when you're up against a deadline.

Overall, the quality and quantity of this collection makes it a worthwhile purchase for LightWave users who need building surfaces and would rather not spend the time and effort to create their own.

Figure 1: SurfacesROM matched, tiled maps for stone wall, birch tree, marble tile, dried mud, and pink marble applied to objects in LightWave.

Dosch Textures vol. 1

reviewed by Doug Kelly

US $89
Dosch Design
Gotthard-Schuell-Str. 14
97828 Marktheidenfeld
Germany
Tel ++49-(0)-9391-915853
http://www.doschdesign.de
info@doschdesign.de

At first glance, the statistics on this collection seem to make it comparable to the SurfacesROM collection. 150 textures, multiple layers, tileable, all at least 1024x1024 and some even 2Kx2K. However, the statistics are less impressive when you look at the subject matter.

These images were apparently rendered with Dosch Designs' DirtyUp shader plug-in, and it shows. The thumbnail catalog closely resembles the test renders that DirtyUp can produce from default settings. The results are images that would be acceptable at 512x512 or even lower resolution; the extra resolution is simply wasted space. The utility of the images is questionable: how many different versions of plaster, rust or slime molds do you need? If you need all those different versions, don't you also need fine control and the ability to animate them?

An old saying goes, "Why keep a cow when milk is so cheap?" This "milk" is overpriced. If you need images like these, I recommend you buy the "cow" - Dosch Designs' DirtyUp plug-in - and render your own texture maps.

GeForce 256

reviewed by Julian Kain

NVIDIA GeForce 256
NVIDIA Corporation
408.615.2500
http://www.nvidia.com
Estimated street price:
$250 to $300
(Note: DDR versions will be released shortly for approximately $100 more than standard SDRAM models.)

Welcome once again to *Keyframe's* graphics chip roundup, where sub-$500 graphics chips are put to the test for LightWave performance. Though I began my reviews for *Keyframe* with the intention of testing inexpensive, consumer level graphics chips under LightWave, it seems that the very definition of "consumer level" is now determined only by price. In the past few years, computer graphics hardware has seen rapid development, strong competition, and falling prices. However, every once in a while, something so revolutionary comes along that it immediately marks the obsolescence of all previous products. This time has come again, and the messiah of the hour is NVIDIA's GeForce 256 graphics chip. Boasting scores of up to five times the speed of our previous champion (3Dlabs Oxygen VX1), the sub-$300 GeForce-based cards can do everything, and they do it very, very well.

The test system is a Pentium II 400MHz with 256MB RAM, using 16-bit color (unless otherwise mentioned). I am currently using 1024x768 as a benchmarking resolution, and would like to hear your feedback on this choice. If I find the majority of artists use a higher resolution, I will gladly adjust the benchmarks. The "Hummer Test", the main benchmark measuring LightWave OpenGL performance, is conducted by doing the following: Load the Hummer scene . from the directory "\NewTek\Vehicles\Hummer" or the LightWave CD. In Layout Options, set the bounding box threshold to 56585. Also, make sure OpenGL is selected for Layout View Graphics. Find a watch or clock that enables you to keep time to the second, and click on the play button at the bottom of the screen. Allow the scene to play through frame 50 and stop timing. Note that your CPU does play a role in the speed you will see. However, increases and decreases in speed will be proportional among graphics chips, making this a valid benchmark. The Hummer Test relies heavily on your OpenGL hardware and is geared mainly towards measuring polygon speed. Other benchmarks in this article include the X29 and the Viewperf 6.1.1 suite. LightWave 5.6c is tested under both Windows NT 4.0 (SP6) and Windows 98 SE. An addition of a LightWave 6.0 test under Windows 2000 is expected soon, possibly by the next issue.

Hummer play time: 25 seconds in Windows NT, 27 seconds in Windows 98

Since shortly after the release of NVIDIA's popular TNT chip (see *Keyframe* #12), hype had begun building around their the highly secret "NV10". Internet rumors were plentiful, proclaiming a chip that would "have more transistors than a Pentium III CPU", and "be able to perform all graphics instructions on the chip". The NV10 was officially named the GeForce 256, and the rumors are true. The most striking thing about the highly anticipated GeForce is the fact that it is a GPU, or graphics processing unit, not your standard 3D chip. Until the release of the GeForce, base graphics instructions were performed by the system CPU and sent to the card for rendering, in both 2D and 3D. With the introduction of the GPU, this bandwidth-hogging system is no more. While the CPU is still used for mathematical calculations in the program or game you're using (such as calculating bone displacement in LightWave or artificial intelligence in games), graphic calculations will no longer have to share CPU time. The result is pure speed. The CPU speed becomes much less of a factor in the graphical output, and produces similar results. Take the Hummer test scene, for example. For each frame, the computer must calculate movement of many objects, movement of the camera, and the resulting changes in lighting. Traditionally, the movement of items and graphical calculations were processed by the CPU, sent through the graphics chip, and out to the monitor. However, with the GeForce, a GPU in itself, all graphical processing (including lighting and the display of all 56584 polygons) is calculated independently of the CPU, and the data sets are combined. Because of independent processing, what we see are results ranging from 2x to 4x the speed of even the fastest and most recent graphics chips available.

The GeForce 256 is a very technically impressive chip, boasting about 23 million transistors, more than AMD Athlon or Intel Coppermine CPUs. It also has an extremely high fill rate (the rate at which polygons are drawn and filled) of 480 Megapixels, as opposed to the standard of less than 366 Megapixels in most current graphics chips. This directly affects polygon speed, LightWave OpenGL performance and your productivity. The GeForce also uses 4 parallel rendering pipelines to prevent speed bottlenecks. Most video cards making use of the GeForce (including the Creative Labs 3D Blaster Annihilator that was tested for this review) use 32MB RAM, though the chip supports the truly massive amount of 128MB. This memory comes in two formats, SDR-RAM (single data rate) and DDR-RAM (double data rate). As of press time, only SDR models were commercially available; DDR models should be arriving by the time you read this, at approximately $100 more than SDR models. Benefits of DDR will be discussed later in this review.

Upon installing the GeForce and NVIDIA drivers (Detonator version 3.53), positive changes were immediately noticeable. Scrolling in any given window was smoother than ever before, and colors were as vivid as the Oxygen VX1. In Windows 98, the Creative Labs drivers caused some DirectX problems (as in failing to run anything at all), but these were immediately resolved with the NVIDIA drivers and a re-installation of DirectX 7. In Windows NT, installation was flawless, and the result was a joy to behold. Flicker-free resolutions of up to 2048 x 1536 in true color are said to be supported, though I was unable to test this due to monitor limitations. Nevertheless, colors were beautiful and text was consistently sharp. Games in both DirectX and OpenGL were astoundingly fast, and increases in resolution rarely caused decreases in speed.

OpenGL speed in LightWave was amazing, to say the least. The GeForce came extremely close to doubling the speed of the previous champion, the Oxygen VX1, whereas the VX1 was only six percent faster than the TNT. Polygon shading was excellent, and colors were once again very vivid. Because of the GeForce's ability to handle a large number of on-screen polygons, adding several more models to the scene barely slowed the GeForce at all. There were some noticeable problems in LightWave, however. There was much more flickering noticeable than with previous graphics chips, and the track created by the Hummer consistently became invisible about once every two frames. Because I didn't see this in other applications or games, I assume it's a fault with LightWave's OpenGL programming, though I can't be sure. Modeler performance was also excellent, and rotating high polygon models was as smooth as rotating a box. Note that high-polygon models are benchmarked in

Viewperf's Advanced Visualizer, where the GeForce score is more than five times the speed of the VX1's. Another plus is that there were no instances of NVIDIA's common "OpenGL hanging" problems, which were present in both the TNT and TNT2.

Whether or not you work with high-polygon models, the GeForce 256 GPU from NVIDIA will significantly increase your OpenGL speed in LightWave. The Windows NT drivers are rock solid, showing a great deal of OpenGL programming effort by the NVIDIA team, and the Windows 98 drivers are nearly as good. Because the chip is a GPU and works independently of the CPU, you will see major speed increases regardless of your CPU model or speed, though the CPU still plays a part in overall LightWave speed, as significant mathematical computations are necessary for animation. Because of the moderate price tag, the only reason I could justify not buying one would be to wait for the DDR model, which allows for faster frame rates at higher resolutions due to enhanced memory speed. Either way, the GeForce is no doubt the first in line of a new breed in GPUs, bringing astounding graphical performance at moderate prices. Even if you consider this kind of speed to be overkill, LightWave 6 will make use of enhanced OpenGL features such as fog and lens flares that call for the kind of performance the GeForce provides.

My rating: 10/10

It's an amazing time in graphics, and NVIDIA has once again pioneered a new direction and produced another A+ product, the GeForce 256 GPU.

Julian Kain is an 18-year-old LightWave hobbyist who hopes to attend Ringling School of Art and Design, which is legendary for its computer animation program, beginning next year. You may contact him via e-mail at julian@juliankain.com and visit his web site and online gallery at http://www.juliankain.com

Benchmark Results Table
(Results for the TNT have been updated, using the new Detonator 3.53 drivers)

Hummer scene:	(lower is better)	Windows 98	Windows NT
	NVIDIA GeForce 256	27 sec.	25 sec.
	Oxygen VX1 (issue #14)	53 sec.	49 sec.
	NVIDIA TNT (issue #12)	50 sec.	50 sec.
X29:	(higher is better)	Windows 98	Windows NT
	NVIDIA GeForce 256	917	1125
	Oxygen VX1 (issue #14)	197	277
	NVIDIA TNT (issue #12)	278	254
Viewperf:	(higher is better)	Windows 98	Windows NT
	NVIDIA GeForce 256		
	ProCDRS	N/A	8.0
	Data Explorer	N/A	24.6
	DesignReview	N/A	19.3
	Advanced Visualizer	N/A	63.0
	Lightscape VS	N/A	2.8
	Oxygen VX1 (issue #14)		
	ProCDRS	N/A	5.9
	Data Explorer	N/A	9.4
	DesignReview	N/A	5.4
	Advanced Visualizer	N/A	12.2
	Lightscape VS	N/A	1.1
	NVIDIA TNT (issue #12)		
	ProCDRS	N/A	9.4
	Data Explorer	N/A	9.5
	DesignReview	N/A	6.2
	Advanced Visualizer	N/A	9.4
	Lightscape VS	N/A	1.2

Animation Traini

NewTek AUTHORIZED TRAINING CENTER

3D Exchange

3D Animation, Modeling, and Effects Training

New Classes • Bigger Facility • Better Location

www.3dexchange.com . (510) 747-1540

Final Cut

by Steve Pugh

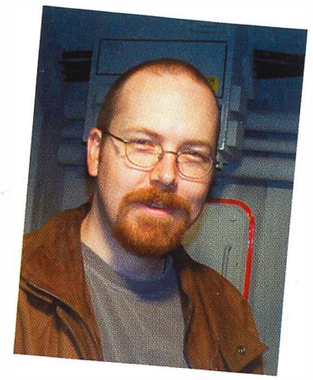

It has long been believed that when we go to that great animation facility in the sky, our lives flash before us in great detail for a brief instant. What if that weren't the case at all, but rather your life was played back for you using whatever technology was available at the time? Sure, maybe people in the old days sort of had to pull up a chair and hope it had good cushions while they listened to a storyteller do his thing. Things picked up with the advent of film and videotape, but you were still a passive observer unless you lucked into some Jimmy Stewart "It's a Wonderful Life" deal.

Now, however, we're at a point where the flashing life can truly be explored! Thanks to DVD technology, your recollection of past triumphs and tragedies is not only crisper and cleaner, but jam-packed with special features previously unavailable to the recently departed! Here's what I'm hoping I get when my time comes.

Chapter Select: If my life is going to flash before my eyes, you can bet there are spots to which I'd like jump back to without wading through the surrounding tedium. Hooking up with Beth Heiserman without having to spend three years processing insurance claims at the same office. Getting my pilot's license without the two years of dreadful Santa Paula weather slowing things down. Scoring a bottle of Vicodin, but without the wisdom tooth extraction. Yes!

Subtitles: Granted, I don't speak any languages other than English, but I bet I would have sounded really cool teaching my dog to sit in Farsi. Who wouldn't?

Alternative Soundtracks and Commentary Tracks: Not that I'd want Philip Glass doing it, thanks. But a good John Williams score would make me feel pretty good about my life, I think. Barring that, Miami Vice's Jan Hammer would surely treat the late '80s nicely. Who would do your commentary? I'd love to have the story of my life recounted by Albert Brooks and porn star Tanya Roberts. I'll bet they'd make a darling couple, too.

Behind the Scenes Footage: Rare glimpses into the mind of the director can be so entertaining. What was Steve thinking when he set fire to a spoonful of flash powder without moving the rest of it out of the way? How did those special effects wizards cut the tip of his thumb off with a cheese slicer? These are the little tricks that people love to learn about.

Deleted Scenes: Finally, see what Steve's real vision was. Why, without the meddling of those studio executives, Steve would have gone to the prom with that girl Mayra. Great fight scenes, like the one in seventh grade which was deleted when Steve ran the other way, would be restored to their glory. Considering the subject matter, though, I'm not sure I'd like to see alternate endings. Too little, too late.

Yes, this new wave of technology is definitely a good thing. Now, as soon as mini-satellite dishes make it into the afterlife, I'll be able to watch two hundred channels of the exact same life, all starting at different times, with no commercial interruption. I can't wait.

Steve Pugh is the General Manager at Foundation Imaging. If he's not daydreaming or fighting a mental block, you can contact him through their website, www.foundation-i.com.

LightWave® [6]

Allow yourself to be seduced.

Seduced by the power, the elegance and the flexibility. Your first encounter will reveal why LightWave dominates the demanding world of show business.

LightWave has always been the industry standard for exceptional realism. Today it just got better. Much better. Now, with caustics, radiosity and integrated Hypervoxels, LightWave [6] understands and delivers the nuances of lighting and volumetrics better than any other animation system on the market. The addition of a 96-bit floating-point render pipeline assures you ultimate color depth and accuracy.

LightWave [6] introduces the world's fastest and most accurate hybrid inverse/forward kinematics engine. Also, unique to LightWave [6] are IntelligEntities™. These include Endomorphs™ and Skelegons™ used to create and control character animation with efficiency and fluidity.

With the release of LightWave [6] nothing was left to chance. We completely redesigned the program and the process with you, the artist, in mind. The joy of LightWave, all the features, all the functions — better, faster, fun.

The seduction is complete.

NewTek booth #535 at Siggraph '99

www.newtek.com
Call **(800) 862-7837**
Outside the U.S. call (210) 370-8000

LightWave 3D is a registered trademark of NewTek, Inc. Hypervoxels, IntelligEntities, Endomorphs and Skelegons are trademarks of NewTek, Inc. © NewTek, 1999. All rights reserved.

NewTek